

PAPER • OPEN ACCESS

Quantum mixed state compiling

To cite this article: Nic Ezzell *et al* 2023 *Quantum Sci. Technol.* **8** 035001

View the [article online](#) for updates and enhancements.

You may also like

- [Variational quantum compiling with double Q-learning](#)

Zhimin He, Lvzhou Li, Shenggen Zheng et al.

- [Large gradients via correlation in random parameterized quantum circuits](#)

Tyler Volkoff and Patrick J Coles

- [Introduction to topological quantum computation with non-Abelian anyons](#)

Bernard Field and Tapio Simula



OPEN ACCESS

RECEIVED
28 September 2022REVISED
17 February 2023ACCEPTED FOR PUBLICATION
16 March 2023PUBLISHED
4 April 2023

Original Content from
this work may be used
under the terms of the
Creative Commons
Attribution 4.0 licence.

Any further distribution
of this work must
maintain attribution to
the author(s) and the title
of the work, journal
citation and DOI.



PAPER

Quantum mixed state compiling

Nic Ezzell^{1,2,*} , Elliott M Ball³ , Aliza U Siddiqui^{4,5} , Mark M Wilde^{4,6} , Andrew T Sornborger^{2,7} , Patrick J Coles^{7,8} and Zoë Holmes^{2,9}

¹ Department of Physics & Astronomy, University of Southern California, Los Angeles, CA 90089, United States of America

² Information Sciences, Los Alamos National Laboratory, Los Alamos, NM 87545, United States of America

³ Physics Department, Lancaster University, Lancaster LA1 4YB, United Kingdom

⁴ Hearne Institute for Theoretical Physics, Department of Physics and Astronomy, and Center for Computation and Technology, Louisiana State University, Baton Rouge, LA 70803, United States of America

⁵ Division of Computer Science and Engineering, Louisiana State University, Baton Rouge, LA 70803, United States of America

⁶ School of Electrical and Computer Engineering, Cornell University, Ithaca, NY 14850, United States of America

⁷ Quantum Science Center, Oak Ridge, TN 37931, United States of America

⁸ Theoretical Division, Los Alamos National Laboratory, Los Alamos, NM 87545, United States of America

⁹ Institute of Physics, Ecole Polytechnique Fédérale de Lausanne (EPFL), CH-1015 Lausanne, Switzerland

* Author to whom any correspondence should be addressed.

E-mail: naezzell@proton.me

Keywords: variational quantum algorithm, quantum mixed states, Hilbert-Schmidt distance, mixed state compiling, quantum principal component analysis, quantum state compression

Abstract

The task of learning a quantum circuit to prepare a given mixed state is a fundamental quantum subroutine. We present a variational quantum algorithm (VQA) to learn mixed states which is suitable for near-term hardware. Our algorithm represents a generalization of previous VQAs that aimed at learning preparation circuits for pure states. We consider two different ansätze for compiling the target state; the first is based on learning a purification of the state and the second on representing it as a convex combination of pure states. In both cases, the resources required to store and manipulate the compiled state grow with the rank of the approximation. Thus, by learning a lower rank approximation of the target state, our algorithm provides a means of compressing a state for more efficient processing. As a byproduct of our algorithm, one effectively learns the principal components of the target state, and hence our algorithm further provides a new method for principal component analysis. We investigate the efficacy of our algorithm through extensive numerical implementations, showing that typical random states and thermal states of many body systems may be learnt this way. Additionally, we demonstrate on quantum hardware how our algorithm can be used to study hardware noise-induced states.

1. Introduction

The task of learning an unknown $d \times d$ quantum state ρ is a fundamental primitive for quantum computing. The well known method of full quantum state tomography learns the matrix elements of ρ directly. The original scheme employs $\mathcal{O}(d^2/\epsilon^2)$ Pauli measurements to learn the state up to additive error ϵ in trace distance [1]. Subsequent refinements require only $\mathcal{O}((\text{rank}(\rho)d/\epsilon^2) \cdot \ln(d/\epsilon))$ measurements [2], or $\mathcal{O}(d/\epsilon^2)$ allowing for a small failure probability [3]. A recent improvement has found it is necessary to use at least $\Omega(\text{rank}(\rho)d/\epsilon)$ measurements, and it was also conjectured to be sufficient [4]. In all cases, tomography aims to obtain a classical description of a quantum state, and as such, the number of measurements it requires scales exponentially with the number of qubits of the target state.

An alternative approach is to give a more operational meaning to learning. In practice, we are often not interested in the exact form of ρ but instead in its properties; i.e. we wish to estimate $o_i := \text{Tr}[\rho O_i]$ for some observable O_i . Aaronson [5] proves that if each observable O_i , for $1 \leq i \leq M$, is restricted to two-outcome measurements, then only $\tilde{\mathcal{O}}(\epsilon^{-4} \log^4 M \log d)$ copies of ρ are sufficient to estimate each o_i up to additive error ϵ by using a method called *shadow tomography* [6–8]. These results were expanded to simpler, experimentally tractable Clifford measurements and random Pauli measurements in [6, 9]. In particular, it was shown that a

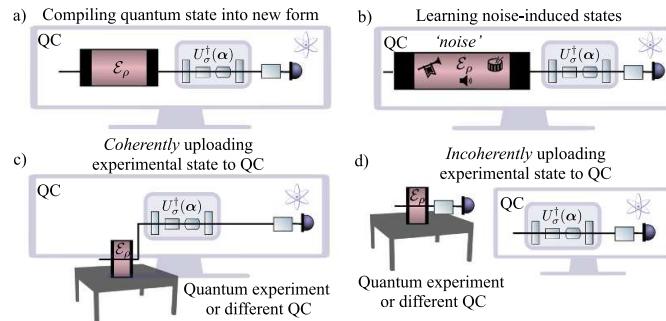


Figure 1. Applications of mixed state compilation. (a) A circuit to prepare the target state is already known and quantum compilation is used to learn a more efficient circuit to prepare that state. (b) The target state is the (unknown) output of (unknown) noise processes on quantum hardware. (c), (d) Quantum compiling is used to upload a quantum state from an experimental system (or different quantum computer) to a quantum computer. This could be performed either coherently using a Loschmidt echo/SWAP test to compute the cost (c) or incoherently by combining our algorithm with classical shadows techniques (d). In all sub-figures \mathcal{E}_ρ denotes the channel that prepares ρ and $U_\sigma^\dagger(\alpha)$ denotes a single parameterized unitary to learn ρ via the CCPS ansatz.

polynomial number of Clifford (Pauli) measurements is sufficient to estimate any low rank (local) observable. The polynomial scaling achieved by shadow tomography is gained at the cost of providing only a partial description of the quantum state.

Yet another approach to quantum state learning, and the one we take here, is instead to solve the problem of state *compilation*. This involves learning a quantum circuit with which we may prepare an approximation of the target state. In this paper, we propose and demonstrate a near-term algorithm for approximately compiling an unknown quantum state by variationally training a PQC. Similar to the shadow tomography case, one can use the final cost function value in our algorithm to place bounds on the deviation of observables estimated from the learnt state from those of the target state.

While the variational compilation of pure states has been explored in [10–12], here we focus on mixed state compilation. Specifically, we propose two different ansätze that can compile a mixed state either into a purification using ancilla or as a convex combination of pure states (CCPSs). The latter is similar to the Hamiltonian-based model approach in [13]. In addition, our approach shares similarities with [14, 15], which present algorithms to learn the diagonalization of a target state, but in contrast to our approach, these methods only apply to low-rank quantum states.

Quantum state compilation, in providing a quantum circuit description of a state, serves a different purpose from tomography or classical shadows, as summarized in figure 1. In the first case, it may be used to ‘upload’ an unknown state of an experimental quantum system to a quantum computer. In other words, given a quantum system in an unknown target mixed state, our algorithms can learn a circuit representation of it for implementation on any digital quantum device. In this case, we can use the compilation to more efficiently store the state for later processing, compute properties of the state that are hard to measure directly, or use the state as the input to another algorithm, say for quantum simulation. Alternatively, a circuit to prepare the target state might already be known, but the aim would be to learn a more efficient (i.e. shorter depth and/or more noise resilient) circuit to prepare that state. These applications are particularly interesting when the unknown mixed state is generated by unknown noise as in (b). In this case, the compilation serves as a snapshot of how a noisy quantum computer corrupted a desired input state. This snapshot can then be used later or perhaps even on a more coherent machine for reliable processing.

In some cases, it may be of interest not to learn the target state perfectly but rather attempt to learn a low(er) rank approximation of it. In general, the resources required to manipulate a quantum state on quantum hardware grow with its rank. Thus, learning a lower rank approximation provides a means of compressing a state to store it more efficiently. The compression of quantum data has a long history tracing back to the early days of quantum information theory [1, 16, 17]. More recently, the question of how well any given state may be approximated using a lower rank state was addressed analytically in [18] for the Hilbert–Schmidt (HS) and trace distances. Interestingly, the optimal lower rank approximation essentially corresponds to performing principal component analysis (PCA) (jump to equation (3) to look ahead) for the HS distance (but not for trace distance). Hence, as we will show, one can also use our learned compilation to perform PCA with a desired cut-off rank. At the same time, the analytical expression for the optimal state from this work provides a natural benchmark for our learning task.

Our algorithm, which is summarized in figure 2, involves variationally minimizing a cost function that is formulated in terms of the HS distance between the target state and ansatz state. This cost can be efficiently

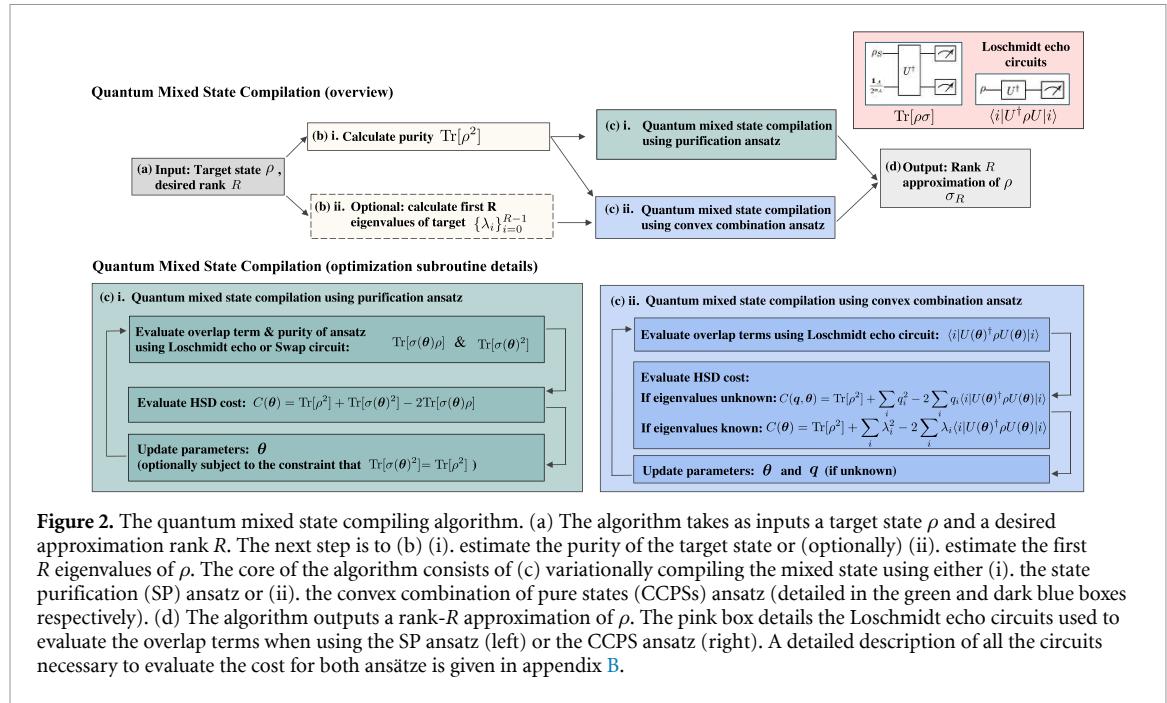


Figure 2. The quantum mixed state compiling algorithm. (a) The algorithm takes as inputs a target state ρ and a desired approximation rank R . The next step is to (b) (i). estimate the purity of the target state or (optionally) (ii). estimate the first R eigenvalues of ρ . The core of the algorithm consists of (c) variationally compiling the mixed state using either (i). the state purification (SP) ansatz or (ii). the convex combination of pure states (CCPSs) ansatz (detailed in the green and dark blue boxes respectively). (d) The algorithm outputs a rank- R approximation of ρ . The pink box details the Loschmidt echo circuits used to evaluate the overlap terms when using the SP ansatz (left) or the CCPS ansatz (right). A detailed description of all the circuits necessary to evaluate the cost for both ansätze is given in appendix B.

measured using either a SWAP test or a Loschmidt echo circuit. In contrast to [19], which sketches a quantum generative adversarial neural network that might be used to learn a mixed state, we further present local variants of our cost functions to mitigate the trainability barrier posed by barren plateaus [20–30].

In the numerical simulations of our proposed algorithm, we demonstrate the applicability of both the purification and convex combination ansätze for learning full and lower rank approximations of a target mixed state. In particular, we numerically simulate the learning of typical random states, as well as of random thermal states of the Heisenberg XY model. We additionally implement our algorithm on quantum hardware in order to compile an unknown state generated by hardware noise.

2. Quantum mixed state compiling (QMSC) algorithm

The QMSC algorithm takes as input a mixed state ρ and a desired approximation rank R . Though not necessary, we assume $R \leq \text{rank}(\rho) =: r$ in defining our algorithm since this is both sensible and simplifies our discussion. In our algorithm, we take this a step further and constrain $R \leq r_\epsilon$, where r_ϵ is the ϵ -rank¹⁰ of ρ , which counts the number of eigenvalues greater than $\epsilon > 0$. This notion of approximate rank is designed to capture the intuition that the contribution of very small but non-zero eigenvalues can often in practice be ignored. The goal then is to optimize the classical parameters α of a parameterized trial state $\sigma(\alpha, R)$ satisfying $\text{rank}(\sigma(\alpha^*, R)) = R$, such that, for the optimized parameters α^* , the output state $\sigma(\alpha^*, R)$ well approximates the target ρ .

In order to assess the closeness between our training state σ and our target state ρ , we employ the following cost function in terms of the HS distance between the two states:

$$\begin{aligned} C(\alpha, R) &\equiv \|\rho - \sigma(\alpha, R)\|_2^2 \\ &:= \text{Tr}[\rho^2] + \text{Tr}[\sigma(\alpha, R)^2] - 2\text{Tr}[\rho\sigma(\alpha, R)]. \end{aligned} \quad (1)$$

As discussed in more detail in sections 2.3 and 2.4, the HS distance can be efficiently computed using a combination of SWAP tests and/or a Loschmidt-echo-like circuit. In appendix A we detail how this cost can be reformulated such that it requires only local measurements [21] to mitigate the barrier to trainability posed by barren plateaus [20–30].

If trained well, the learned state $\sigma(\alpha^*, R)$ should be close to the solution to the quantum low-rank approximation problem (QLRAP) [18]. As discussed in [18], the unique optimal state that minimizes the HS distance, subject to a rank constraint, i.e. the state $\sigma(\alpha_{\text{opt}}, R)$ where

$$\alpha_{\text{opt}} := \arg \min C(\alpha, R), \quad (2)$$

¹⁰ We note that the notion of ϵ -rank was also used in [31], but its formal definition is slightly different than the one we choose to use here. A smooth version of the max-entropy (or $\alpha=0$ Renyi entropy) provides yet another definition for ϵ -rank that is distinct from ours [32, 33].

takes the form

$$\sigma(\boldsymbol{\alpha}_{\text{opt}}, R) = \tau_R + \left(\frac{1 - \text{Tr}[\tau_R]}{R} \right) \Pi_R, \quad (3)$$

$$\tau_R := \Pi_R \rho \Pi_R. \quad (4)$$

Here Π_R is a projector onto the eigenstates corresponding to the R largest eigenvalues of ρ . That is, Π_R is a projector onto the first R principal components of ρ . We note here that the projection Π_R is the same conceptually as the typical subspace projection [34] used in quantum data compression [16] (given that it projects onto the high probability subspace of the state ρ), and the quantum channel in (3) (i.e. $(\cdot) \rightarrow \Pi_R(\cdot)\Pi_R + \text{Tr}[(I - \Pi_R)(\cdot)]\frac{\Pi_R}{R}$) is essentially the same as the encoding channel used in quantum data compression [1, equation (12.50)].

With this analytical solution, we can define a natural performance metric for the QMSC algorithm as the difference between the found cost and optimal possible cost,

$$\Delta_R := C(\boldsymbol{\alpha}^*, R) - C(\boldsymbol{\alpha}_{\text{opt}}, R). \quad (5)$$

More explicitly, this expression takes the form

$$\Delta_R = C(\boldsymbol{\alpha}^*, R) - \sum_{i=R}^{r-1} \lambda_i^2 + \sum_{i=0}^{R-1} (\lambda_{\sigma i} - \lambda_i)^2, \quad (6)$$

where $\{\lambda_i\}_{i=0}^{r-1}$ denotes the set of eigenvalues of the target state. Of course when ρ is truly unknown, we can only report the optimal cost $C(\boldsymbol{\alpha}^*, R)$, but in this work, we actually compute Δ_R to verify that our algorithm is working as intended.

In this paper, we develop a practical, noisy intermediate-scale quantum (NISQ)-friendly algorithm to find $\sigma(\boldsymbol{\alpha}^*, R)$, an approximation of $\sigma(\boldsymbol{\alpha}_{\text{opt}}, R)$. In doing so, we also find the first R principal values and components of ρ , and so we are essentially performing quantum PCA, analogous to some previously proposed algorithms [14, 15, 35]. The precise form in which $\sigma(\boldsymbol{\alpha}^*, R)$ (or the first R principal components) is obtained depends on our choice of ansatz. We consider two different ansätze, one based on learning an approximate purification of ρ and the second on decomposing an approximation of it into a CCPSSs. We also discuss the complexity of computing our cost function, as well as its operational meaning. These are detailed in the following sections. A summary of our algorithm, which also describes the two ansätze choices, is given in figure 2.

2.1. Complexity of cost function

For all variational quantum algorithms (VQAs), the purpose of using the quantum computer in the optimization loop is to estimate a cost function efficiently, which otherwise would be difficult to estimate classically. Hence, it is helpful to establish that the cost function is indeed classically hard to estimate. Previously proposed VQAs for quantum compiling [10] and linear system solving [36] have established classical hardness via the DQC1 hardness of estimating the relevant cost function. (In the previous sentence, DQC1 stands for deterministic quantum computation with one clean qubit [37]). We can make this same argument for the cost function in equation (1), as follows.

A special case of computing equation (1) is when the two states happen to be Choi states associated with unitary processes acting on a d dimensional Hilbert space. In this case, we write $\rho = |\phi_U\rangle\langle\phi_U|$ and $\sigma = |\phi_V\rangle\langle\phi_V|$, where $|\phi_U\rangle = (\mathbb{1} \otimes U)|\phi\rangle$, $|\phi_V\rangle = (\mathbb{1} \otimes V)|\phi\rangle$, and $|\phi\rangle = (1/\sqrt{d})\sum_j |j\rangle|j\rangle$ is the standard Bell state. Then the cost function in (1) becomes $C = 2 - 2\text{Tr}(\rho\sigma)$ where $\text{Tr}(\rho\sigma) = |\text{Tr}(U^\dagger V)|^2/d^2$. Hence, in this special case we have that $C = 2C_{\text{HST}}$ where C_{HST} is the HS test cost of [10]. Estimating the latter was shown to be DQC1-hard in [10], and hence our cost function C is also DQC1-hard. Since efficient classical simulation of DQC1 would imply a collapse of the polynomial hierarchy [38, 39], standard complexity assumptions imply the classical hardness of estimating C .

2.2. Operational meaning of cost function

2.2.1. Observable estimation

There is a close connection between the HS distance and trace distance for low-rank states [40], which gives operational meaning to our cost function in this case. Namely when at least one of the states is low rank, then the two distance measures are essentially equivalent [40]. This can be seen from the following inequality relating the 2-norm to the 1-norm:

$$\|\rho - \sigma\|_2^2 \leq \|\rho - \sigma\|_1^2 \leq 4\mathcal{R} \cdot \|\rho - \sigma\|_2^2 \quad (7)$$

for any two density matrices ρ and σ . Here, $\mathcal{R} = \text{rank}(\rho)\text{rank}(\sigma)/[\text{rank}(\rho) + \text{rank}(\sigma)]$ is a quantity called the reduced rank, which is analogous to the reduced mass employed in physics.

In this sense, one can use the HS distance, and hence our cost function, as a strong upper bound on the trace distance. In turn, the trace distance has operational meaning in terms of the difference of observable expectation values on the two states [1], $\|\rho - \sigma\|_1 = 2 \max_P \text{Tr}(P(\rho - \sigma))$ where the maximization is over all POVM (positive operator valued measure) elements P , i.e. operators satisfying $0 \leq P \leq \mathbb{1}$. Hence, our cost function inherits this operational meaning:

$$C(\alpha, R) \geq \frac{1}{\mathcal{R}} [\text{Tr}(P\rho) - \text{Tr}(P\sigma(\alpha, R))]^2 \quad (8)$$

for any POVM element P , where \mathcal{R} is the reduced rank for ρ and $\sigma(\alpha, R)$.

2.2.2. Eigenvalue and eigenvector estimation

We also establish an operational meaning for our cost function in the context of minimizing the errors in the eigenvectors and eigenvalues of the compiled state. This is particularly relevant in the context of using our algorithm for PCA, since PCA precisely aims to extract the eigenvectors with the largest eigenvalues.

First, we note that our proposed cost is an upper bound on the squared difference between the true eigenvalues of ρ and the eigenvalues of the learnt state σ , as follows. Let us denote the eigenvalues of ρ in increasing order as $\{\lambda_i\}_{i=0}^{d-1}$ and the eigenvalues of σ as $\{\mu_i\}_{i=0}^{d-1}$, with d the Hilbert space dimension. Then a measure of eigenvalue error is

$$\Delta_\lambda := \sum_{i=0}^{d-1} (\lambda_i - \mu_i)^2. \quad (9)$$

It then follows from the Hoffman–Wielandt theorem that

$$\Delta_\lambda \leq \|\rho - \sigma\|_2^2 = C. \quad (10)$$

Thus if one can achieve a small cost value, the average error in the learnt eigenvalues is guaranteed to be small.

We remark that a small cost value might not be achievable since the optimal cost $C(\alpha_{\text{opt}}, R)$ is often non-zero. In this case, one could consider an alternative measure of eigenvalue error that, unlike (9), would vanish for the optimal state in (3). Namely, one can consider the set S_R of states that have the correct R -largest eigenvalues up to an additive constant, and then define a measure of eigenvalue error as $\Delta_{S_R}(\sigma) = \min_{\tau \in S_R} \|\sigma - \tau\|_2^2$. Then it is clear that our cost function upper bounds this error as well: $C \geq \Delta_{S_R}$.

Second, we note that our proposed cost function is an upper bound on the eigenvector error measure introduced in [14]. Specifically, a natural measure of the difference between the true eigenvectors of the target state, $\{|v_i\rangle\}_{i=0}^{r-1}$, and the learnt eigenvectors, $\{|u_i\rangle\}_{i=0}^{R-1}$, is given by

$$\Delta_v := \sum_{i=0}^{R-1} \langle \delta_i | \delta_i \rangle \quad (11)$$

where

$$|\delta_i\rangle := \rho|u_i\rangle - \mu_i|u_i\rangle. \quad (12)$$

Expanding this out we have that

$$\Delta_v = \sum_{i=0}^{R-1} (\langle u_i | \rho^2 | u_i \rangle + \mu_i^2 - 2\mu_i \langle u_i | \rho | u_i \rangle) \leq \|\rho - \sigma\|_2^2 \quad (13)$$

where for the inequality we use the fact that

$$\sum_{i=0}^{R-1} \langle u_i | \rho^2 | u_i \rangle \leq \sum_{i=0}^{d-1} \langle u_i | \rho^2 | u_i \rangle = \text{Tr}[\rho^2], \quad (14)$$

(i.e. the trace can be taken with respect to an arbitrary orthonormal basis). Thus a small cost function guarantees a small eigenvector error by this measure.

2.3. State purification (SP) ansatz

Let us now move onto the different types of ansatz constructions. The SP ansatz constructs trial states of the form

$$\sigma_{\text{SP}}(\boldsymbol{\theta}, n_A) := \text{Tr}_A[U_{\boldsymbol{\theta}}(|0\rangle\langle 0|)^{\otimes(n+n_A)}U_{\boldsymbol{\theta}}^\dagger], \quad (15)$$

where $U_{\boldsymbol{\theta}}$ acts on the n system qubits (the same qubits where the state ρ resides) plus an ancilla register A composed of n_A qubits. In using n_A ancillas, we guarantee that $\text{rank}(\sigma) \leq 2^{n_A}$, but for a typical parameterized circuit $U_{\boldsymbol{\theta}}$ with randomly generated parameters, one has precisely $\text{rank}(\sigma) = 2^{n_A}$. Thus, in practice, the SP ansatz is only compatible with controlling R in powers of two, so that $R = 2^{n_A}$. While there are many different possible choices in the ansatz for $U_{\boldsymbol{\theta}}$, we suggest using either a problem-inspired approach obeying the symmetries of the target state [41–43] and/or an adaptive approach [44] to mitigate the problem posed by barren plateaus. A discussion of specific choices in our work for different ensembles is given in appendix B.

Upon substitution of the SP ansatz into our proposed HS distance-based cost function, we obtain

$$C_{\text{SP}}(\boldsymbol{\theta}, n_A) = \text{Tr}[\rho^2] + \text{Tr}[\sigma(\boldsymbol{\theta}, n_A)^2] - 2\text{Tr}[\rho\sigma(\boldsymbol{\theta}, n_A)], \quad (16)$$

where we use S to denote the system of the states ρ and $\sigma(\boldsymbol{\theta}, n_A)$. Below we describe efficient ways of computing each of these terms.

2.3.1. Purity of target

The purity $\text{Tr}[\rho^2]$ of the target ρ may be measured using a SWAP test [45] or its destructive variant [46]. For the estimate to be within ε additive error of the true value with probability not smaller than $1 - \delta$, the Hoeffding bound implies that it suffices to take $O(\varepsilon^{-2} \log \delta^{-1})$ samples. Alternatively, as this term remains constant throughout the procedure, we may opt to neglect it and only focus on optimizing the remaining two parameter-dependent terms.

2.3.2. Purity of ansatz

The purity $\text{Tr}[\sigma(\boldsymbol{\theta}, n_A)^2]$ of the ansatz $\sigma(\boldsymbol{\theta}, n_A)$ may again be measured via a SWAP test. It is also possible to make use of the fact that we are representing it via its purification and measure the purity of the ansatz via the Loschmidt-echo circuit shown in figure 2. Again, from the Hoeffding bound, it suffices to take $O(\varepsilon^{-2} \log \delta^{-1})$ samples.

2.3.3. Overlap term

The overlap term $\text{Tr}[\rho\sigma(\boldsymbol{\theta}, n_A)]$ can be evaluated using either a SWAP test or the Loschmidt-echo type circuit pictured in figure 2. The latter requires fewer qubits and controlled unitaries and hence is more NISQ friendly. To see why this circuit works, note that since we prepare the trial state via a higher-dimensional purification found by evolving an initial all-zero state to the purification of ρ , we can rewrite the overlap term as follows

$$\begin{aligned} \text{Tr}[\rho\sigma] &= \text{Tr}_S[\rho \text{Tr}_A[U_{\boldsymbol{\theta}}(|0\rangle\langle 0|)^{\otimes(n+n_A)}U_{\boldsymbol{\theta}}^\dagger]] \\ &= \text{Tr}_{SA}[(\rho \otimes \mathbb{1}_A)(U_{\boldsymbol{\theta}}(|0\rangle\langle 0|)^{\otimes(n+n_A)}U_{\boldsymbol{\theta}}^\dagger)] \\ &= 2^{n_A} \text{Tr}_{SA} \left[U_{\boldsymbol{\theta}}^\dagger \left(\rho \otimes \frac{\mathbb{1}_A}{2^{n_A}} \right) U_{\boldsymbol{\theta}}(|0\rangle\langle 0|)^{\otimes(n+n_A)} \right]. \end{aligned} \quad (17)$$

This corresponds to preparing a maximally mixed state on the n_A qubit ancillary system alongside ρ which is prepared on the system register S , evolving the system and ancilla registers under $U_{\boldsymbol{\theta}}$, and then performing an all-zero measurement. Due to the factor of 2^{n_A} , the number of shots required to measure this cost within additive error ε will scale exponentially in the number of ancilla qubits n_A . Hence the Loschmidt echo method for computing the overlap is only appropriate for learning low rank approximations to ρ where $n_A \in \text{poly}(\log(n))$.

As discussed earlier, the above method may also be used to calculate the purity of our ansatz simply by replacing the initial state ρ with the trial state σ . However, the purity of our target may not be measured this way, as it requires knowledge of a purification of ρ , which is precisely what we are trying to find.

2.4. CCPSs ansatz

The CCPSs ansatz constructs trial states of the form

$$\sigma_{\text{CCPS}}(\boldsymbol{\alpha}, R) := \sum_{i=0}^{R-1} p_{\boldsymbol{\phi}}(i) U_{\boldsymbol{\theta}} |i\rangle\langle i| U_{\boldsymbol{\theta}}^\dagger. \quad (18)$$

Here $\{|i\rangle\}_{i=0}^{R-1}$ denotes a subset of the computational basis of n qubits, $\alpha = (\theta, \phi)$ is a vector of parameters, U_θ is a parameterized quantum circuit (PQC), and p_ϕ is a parameterized probability distribution. An appealing feature of this ansatz is that learning a rank- R approximation gives any rank $R' < R$ approximation for free since one can always ‘drop’ the eigenstates corresponding to the smallest $R - R'$ eigenvalues and then re-normalize (see equation (31a)).

Since the state that minimizes our cost, $\sigma(\alpha_{\text{opt}}, R)$ (i.e. the solution to the QLRAP given in equation (3)), is proportional to the first R principal components of the target state ρ , it follows that this ansatz can be used to learn the principal components of ρ . More precisely, we have that the first R principal values of ρ are given by $\{p_\phi(i)\}_{i=0}^{R-1}$ and its principal components are $\{U_{\theta_{\text{opt}}} |i\rangle\}_{i=0}^{R-1}$.

When the rank of the trial state is low, i.e. $R \in \Omega(\text{poly}(n))$, such that we are learning a low rank approximation of the target state ρ , or for small scale problems, the vector of probabilities $\mathbf{p}_\phi := (p_\phi(0), \dots, p_\phi(R-1))$ may be stored simply as a classical vector. For learning high rank approximations to larger full rank states, the probability vector is exponentially large and thus cannot be explicitly stored efficiently. Rather the process will need to be sampled from. Such samples can be generated via classical neural networks, such as generative neural networks [47] or Boltzmann machines [48]. In this case, there are similarities between the CCPS ansatz and the Hamiltonian models considered in [13]. Similarly to the SPA, for U_θ we suggest using either a problem inspired ansatz obeying the symmetries of the target state [41–43] and/or an adaptive approach [44]. A discussion of specific choices in our work for different ensembles is given in appendix B.

Upon substituting the expression for the CCPSs ansatz into our proposed HS distance cost function, we obtain

$$C_{\text{CCPS}}(\alpha, R) = \text{Tr}[\rho^2] + \sum_i p_\phi(i)^2 - 2 \sum_i p_\phi(i) \langle i | U_\theta^\dagger \rho U_\theta | i \rangle, \quad (19)$$

with $\alpha = (\theta, \phi)$. Here we describe how each of the terms in the above cost may be efficiently computed.

2.4.1. Purity of target

As for the SP ansatz, the first term, i.e. the purity of the target state, may be computed using a SWAP test [45] or its destructive variant [46] (see also [49]).

2.4.2. Purity of ansatz

The second term, the purity of the guessed state, is equal to the sum of the square of the probabilities of the parameterized distribution, $p_\phi(i)$. For low rank approximations, i.e. $R \in \text{poly}(n)$, one may store the probability vector classically and thus this term can be computed by basic arithmetic.

When the probability vector is too large to be stored explicitly, but rather is handled via sampling, the purity of the ansatz can be estimated with a classical version of the SWAP test. Indeed, the approach is to take two independent samples from the distribution $p_\phi(i)$ (call the random samples I and J). We then set an indicator random variable χ_P as $\chi_P = 0$ if the samples are not equal, and $\chi_P = 1$ if the samples are equal. The expectation of this random variable is then given by

$$\mathbb{E}[\chi_P] = 0 \cdot \Pr[I \neq J] + 1 \cdot \Pr[I = J] \quad (20)$$

$$= \Pr[I = J] \quad (21)$$

$$= \sum_{i,j} p_\phi(i)p_\phi(j)\delta_{i,j} \quad (22)$$

$$= \sum_i (p_\phi(i))^2. \quad (23)$$

Thus, the random variable χ_P is an unbiased estimator of the collision probability, $\sum_i (p_\phi(i))^2$, and takes values between zero and one. The Hoeffding bound then applies, and we can take $O(\varepsilon^{-2} \log \delta^{-1})$ independent samples of χ_P in order to estimate the collision probability $\sum_i (p_\phi(i))^2$ to within additive error ε with probability not smaller than $1 - \delta$.

2.4.3. Overlap term

A naive approach to computing the third term, the overlap between the guess σ and target ρ , would be to estimate each of the $\langle i | U_\theta^\dagger \rho U_\theta | i \rangle$ terms using the Loschmidt echo circuit shown in figure 2, followed by classical post processing. That is, we compute $\langle i | U_\theta^\dagger \rho U_\theta | i \rangle$ by preparing the state ρ , performing the unitary U_θ^\dagger , and then measuring in the computational basis. The total overlap term could simply be computed by weighting each of the probabilities $\langle i | U_\theta^\dagger \rho U_\theta | i \rangle$ by the corresponding classical probability $p_\phi(i)$ and taking their sum. This naive approach will work for small problems; however, even in the case where $p_\phi(i)$ can be explicitly stored, this method will not be efficient for larger problems. The problem is that one needs enough shots to estimate all 2^n probabilities $\langle i | U_\theta^\dagger \rho U_\theta | i \rangle$. Combining this observation with Hoeffding's equality, it is apparent that this method requires an exponential number of shots, $O(2^n \varepsilon^{-2} \log \delta^{-1})$.

Instead, we propose computing the overlap term using a generalization of the classical SWAP test. Let us define the distribution

$$q_\theta(i) := \langle i | U_\theta^\dagger \rho U_\theta | i \rangle, \quad (24)$$

which can be sampled from via the Loschmidt echo circuit. Then we see that

$$\text{Tr}[\rho\sigma] = \sum_i p_\phi(i) q_\theta(i). \quad (25)$$

This quantity can be estimated by taking a sample from $p_\phi(i)$ and an independent sample from $q_\theta(i)$ (call the samples I and J) and setting an indicator random variable χ_O if the samples are not equal and $\chi_O = 1$ if the samples are equal. The expectation of this random variable is given by

$$\mathbb{E}[\chi_O] = 0 \cdot \Pr[I \neq J] + 1 \cdot \Pr[I = J] \quad (26)$$

$$= \Pr[I = J] \quad (27)$$

$$= \sum_{i,j} p_\phi(i) q_\theta(j) \delta_{i,j} \quad (28)$$

$$= \sum_i p_\phi(i) q_\theta(i). \quad (29)$$

Thus, the random variable χ_O is an unbiased estimator of the collision probability $\text{Tr}[\rho\sigma]$, and it takes values between zero and one. The Hoeffding bound then applies, and we can take $O(\varepsilon^{-2} \log \delta^{-1})$ independent samples of χ_O in order to estimate this collision probability to within additive error ε with probability not smaller than $1 - \delta$.

2.5. Performing PCA with our ansätze

Both ansätze also allow for PCA of ρ . To make this precise, we first write the target state as

$$\rho = \sum_{i=1}^r \lambda_i |v_i\rangle\langle v_i| \quad (30)$$

where $\lambda_1 > \lambda_2 > \dots > \lambda_r$ are the ordered principal values of ρ with associated principal components $|v_i\rangle$.

By inspection, the CCPS ansatz is directly an ansatz for the principal components of ρ . In other words, learning a CCPS representation of $\sigma(\alpha^*, R)$ provides an estimate of the principal components, $|u_i\rangle := U(\theta^*)|i\rangle$, explicitly. It also provides an estimate of the principal values, $p_\phi(i)$, but due to normalization, these are expected to be different from λ_i by an additive constant when Δ_R is small. As R approaches r_ϵ , this additive constant goes to zero, and here $p_\phi(i) \approx \lambda_i$. Nevertheless, even when $R < r_\epsilon$, the values of $p_\phi(i)$ are such that $\sigma(\alpha^*, R)$ acts as the closest rank R proxy, so in this operational sense, it is still appropriate to call $p_\phi(i)$ the principal value estimates.

One very useful property of these explicit principal component/value estimates is that they allow us to construct any numerically optimal $R' < R$ approximation by truncation,

$$\sigma(\boldsymbol{\alpha}^*, R') = \sum_{i=0}^{R'-1} \tilde{p}_\phi(i) U(\boldsymbol{\theta}^*) |i\rangle\langle i| U^\dagger(\boldsymbol{\theta}^*), \quad (31a)$$

$$\tilde{p}_\phi(i) = p_\phi(i) + \left(1 - \sum_{i=0}^{R'-1} p_\phi(i)\right) / R'. \quad (31b)$$

Indeed, this is perhaps one convincing way to view each $p_\phi(i)$ as the appropriate operational sense of ‘principal value’ when we truncate the rank.

Though less obvious, the SP ansatz can also be used for PCA by using a carefully designed circuit for the purification ansatz as described in appendix B. The result is that computational basis measurements on the ancilla system prepare the principal vectors on the target system with probabilities given by the principal values. Thus, the knowledge of principal values/components here is implicit, and hence cannot be directly used to obtain lower rank approximations by truncation.

Finally, we remark that while PCA is a general procedure, it has an intuitive meaning for physically relevant classes of states. As an example, consider an XY thermal state. In general, the eigenvalues follow a Boltzmann distribution, and the eigenvectors are those of the XY Hamiltonian itself. At low temperature, we often say a state is *approximately in its ground-state*. More precisely, we mean,

$$\rho \approx \rho_\epsilon \equiv \sum_{i=1}^q \lambda_i |v_i\rangle\langle v_i|, \quad \sum_{i=1}^q \lambda_i = 1 - \epsilon \quad (32)$$

for ϵ small. In our language, we would say the state is approximately rank q , and when $q \ll r$ as is typical for low temperature thermal states, it is approximately low rank¹¹. By performing PCA with a target rank $R = q$, we find approximations of the eigenvectors with low energies $\{|v_i\rangle\}_{i=1}^q$ and corresponding Boltzmann weights $\{\lambda_i\}_{i=1}^q$. Hence, our PCA algorithm can be thought of as a way to learn the Boltzmann weights and a means to prepare low-lying energy eigenvectors of a quantum thermal state which has also been explored in other NISQ friendly works [50].

2.6. Comparison of ansätze

The reliance on an ancillary system to compute the cost function for the SP ansatz naturally increases the resources required for computation. Furthermore, in general, for a typical choice of U_θ , the guess state will have rank $R = 2^{n_A}$. That is, one is limited to ranks of powers of two. In contrast, CCPS both requires no ancilla and allows for fine control over both output ranks. However, the need to learn p_ϕ under the constraint of convexity, in addition to θ , potentially increases the complexity of our optimization subroutine. Thus the choice as to whether to use CCPS or SP will depend in large part on whether classical resources (optimization power) or quantum resources (qubits available) are more constrained.

Of course, the choice as to whether to use the CCPS or SP ansatz may depend not only on the required resources but also the end goal of the subroutine. For example, if the end goal is PCA, this is more readily performed using CCPS since it automatically finds the first R principal components of the target. Alternatively, one can imagine situations where it is desirable to learn the purification of the target. For example, the purification of a state opens up methods for computing entanglement measures between subsystems of that state [51–53], the fidelity between two states (given Ulmann’s theorem) [54, 55], as well as symmetry measures [56].

2.7. Summary of algorithm

The QMSC algorithm is summarized in figure 2. At its core it is composed of the following steps.

1. Start with a target state ρ , and a desired rank R for the compiled state σ_R used to approximate ρ .
2. Choose whether to learn the target using the purification or convex combination ansatz. In general, this decision will depend on the purpose for which the state is being learned and the resources available. A discussion of the circuits one must run in both cases is given in appendix B.
3. Minimize the HS distance cost C , equation (1), using a hybrid quantum–classical optimization loop to find the trained parameters $\boldsymbol{\theta}^*$ (in the case of the SP ansatz) and $(\boldsymbol{\theta}^*, \phi^*)$ (in the case of the CCPS ansatz) that approximately minimize C . If using a gradient based optimizer one can analytically compute the gradient of the cost using the parameter shift rule [57, 58]. For the case of the purity of the ansatz term for the state

¹¹ We comment that this exact notion of low rank actually precisely agrees with the definition of epsilon-rank used in [31]. We employ a slightly different definition already given that is more amenable to NISQ experiments.

purification ansatz, $\text{Tr}[\sigma(\theta, n_A)^2]$, this rule needs modifying to account for correlations. For more details on computing the gradients, see appendix C.

Here we briefly describe two ways in which prior information about the structure of the target state could be used to make the mixed state learning algorithm more efficient.

In the first instance, the purity of the target state may be computed and then we need only consider ‘guess’ states σ_R with the correct purity. That is, one may use the target purity as a constraint during the minimization of the HS distance cost, equation (1). This amounts to maximizing the overlap term $\text{Tr}[\rho\sigma]$ subject to the constraint $\text{Tr}[\sigma^2] = \text{Tr}[\rho^2]$. We note that this constrained optimization may not be compatible with the rank constraint if the desired rank is much lower than the true rank of the target, i.e. $R \ll r$.

Going a step further, one could also simplify the task of learning a rank- R approximation to ρ by first learning the R largest eigenvalues of ρ and then using the mixed state learning algorithm to learn their corresponding eigenvectors. The largest eigenvalues of any state σ can be learnt non-variationally using the quantum algorithm proposed in [49]. Denoting the measured eigenvalues as $\{\lambda_i\}_{i=0}^{R-1}$, the eigenvectors may be learnt by maximizing the overlap term $\sum_i \lambda_i \langle i | U_\theta^\dagger \rho U_\theta | i \rangle$.

As discussed further in appendix A, one theoretical advantage of these two modifications is that they require optimizing only a single overlap term rather than the difference between two purity terms and an overlap term. This overlap term takes the form of a standard variational quantum eigensolver cost and therefore can be readily transformed into a local cost for which we have trainability guarantees [21].

3. Numerical simulations

We begin our numerical simulations discussion with a brief summary of the target states and chosen optimizer. Additional details can be found in our open source code [59] which includes our raw data and the scripts we used to generate it. We then discuss the results of compiling each of the listed states in separate sections.

3.1. Description of states and optimizer

We discuss the performance of our QMSC algorithm for three main types of states:

1. Random states drawn from the Bures measure.
2. Thermal states of an XY chain with random coefficients at both low and high temperature.
3. Noise-induced states generated by simple circuits on NISQ hardware. (We shall call them NISQ states for short.)

In the first two cases, we perform the entire optimization using idealized classical simulations. That is, we evaluate the cost functions using matrix operations on a classical machine with no error model and no shots (i.e. infinite precision). Henceforth, the use of ‘idealized classical simulations’ will continue to have this precise meaning. This serves both as a proof-of-principle as well as a means to estimate an empirical idealized scaling. For the NISQ states, discussed in section 4, we evaluate all cost functions on quantum hardware but use classical computation for the parameter updates, which is the standard VQA approach.

In each case, we consider a low rank approximation by setting $n_A = 1$ (or, equivalently, $R = 2$) or a full (epsilon) rank approximation with $n_A = \lceil \log_2 r_\epsilon \rceil$ (or $R = 2^{\lceil \log_2 r_\epsilon \rceil}$). It is natural to wonder why we choose to use $R = 2^{\lceil \log_2 r_\epsilon \rceil}$ instead of $R = r_\epsilon$. The reason is straightforward: we simply want both the SP and the CCPS ansatz to attempt to learn the same state to the same rank approximation to make a fair comparison, and we can only control the rank of the SP ansatz in powers of two.

We use the gradient-free Powell optimizer [60] provided in the `scipy` [61] optimization library. We find that Powell is more robust and generally outperforms the common `scipy` black box alternatives such as BFGS [62], Nelder-Mead [63], SLSQP [64], and COBYLA [65] for our problems. Of course, the performance could be improved by using advanced VQA optimizers such as SPSA [66] or ICANs [67, 68], but we do not pursue this refinement since we found Powell to give reasonable results.

3.2. Bures random states results

We first study the Bures random state distribution because it is a reasonable sampling distribution when nothing about the quantum state is known [69, 70] (see appendix D.1 for more details). For some intuition, note that one way to generate n qubit Bures random states is by preparing a Haar random state on $n + n$ qubits, applying an n -qubit ‘local Haar random’ unitary on the system qubits, and then tracing out the ancilla. In this sense, learning the purification is similar to learning a Haar random unitary on $2n$ qubits which we know to be intractable for VQAs due to an intrinsic, ansatz-independent vanishing gradient problem [22].

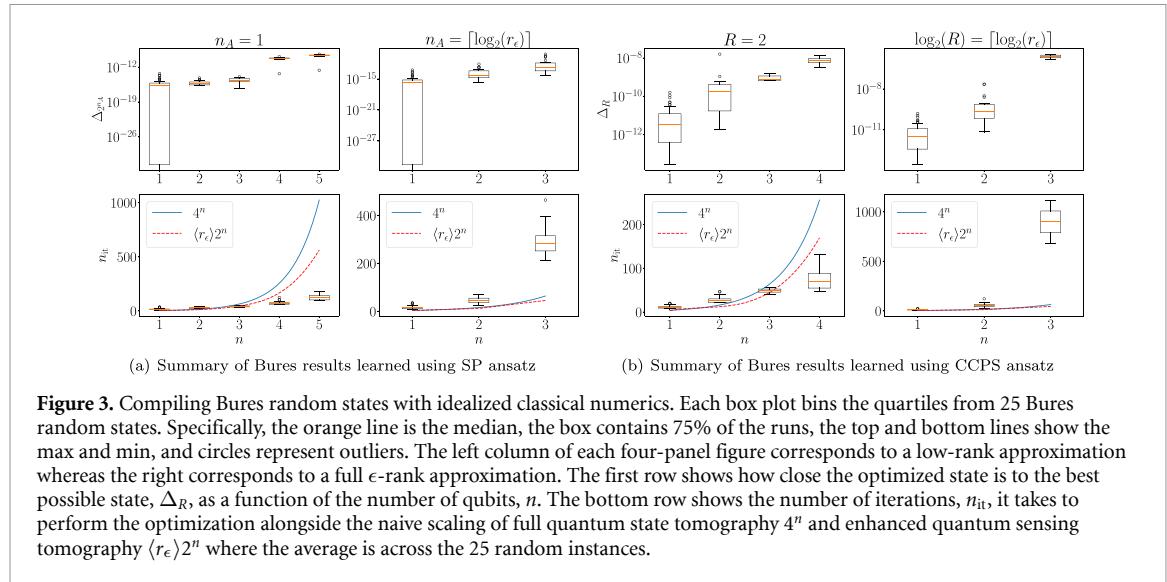


Figure 3. Compiling Bures random states with idealized classical numerics. Each box plot bins the quartiles from 25 Bures random states. Specifically, the orange line is the median, the box contains 75% of the runs, the top and bottom lines show the max and min, and circles represent outliers. The left column of each four-panel figure corresponds to a low-rank approximation whereas the right corresponds to a full ϵ -rank approximation. The first row shows how close the optimized state is to the best possible state, Δ_R , as a function of the number of qubits, n . The bottom row shows the number of iterations, n_{it} , it takes to perform the optimization alongside the naive scaling of full quantum state tomography 4^n and enhanced quantum sensing tomography $\langle r_e \rangle 2^n$ where the average is across the 25 random instances.

Given their lack of structure we see Bures random states as a good test case to compare the SP and CCPS ansätze. However, due to the unavoidable vanishing gradient problem—along with the large number of parameters needed for an unstructured state (see appendix B)—we only test modest sizes. Specifically, we test from $n = 1$ to $n = 4, 5$ for low rank approximations and from $n = 1$ to $n = 3$ for full ϵ -rank approximations (i.e. we learn up to a six-qubit random purification).

We tested our algorithm on 25 Bures random states for each n . The results are shown in figure 3. Here, SP results are shown in (a) and the CCPS results in (b). We plot the difference between the optimized cost and the lowest possible cost Δ_R (see equation (5) or appendix E for more details) as a function of system size, n , for both low rank ($n_A = 1, R = 2$) and full ϵ -rank ($n_A = \lceil \log_2(r_e) \rceil, R = 2^{\lceil \log_2(r_e) \rceil}$) approximations along with the number of iterations n_{it} necessary to reach Δ_R . Here the compilation is performed classically using the Powell optimizer. Our ansatz, which uses alternating layers of arbitrary two-qubit gates, is explained in appendix B.

We focus our discussion first on the top row, where we plot the performance metric Δ_R . Here, we see that our algorithm is capable of learning completely (Bures) random states provided n is small enough. Indeed, for the SP ansatz, the value Δ_R stays below 10^{-10} for all values of n tested when compiling both a low rank and a full ϵ -rank approximation. Interestingly, the values of Δ_R reached for the CCPS ansatz are substantially higher, reaching values of up to 10^{-8} for the low-rank and up to 10^{-6} for the full ϵ -rank approximations. This suggests that the CCPS optimization is more difficult than the SP one. This is plausibly due to the fact that the optimization was performed over both angles and probabilities, which needed to satisfy a normalization constraint.

In the second row of figure 3, we plot the number of iterations, n_{it} , needed to reach the Δ_R values above. For reference, we also plot the curve showing 4^n scaling for naive full tomography as well as a curve $\langle r_e \rangle 2^n$ scaling. Here, $\langle r_e \rangle$ denotes the average ϵ -rank across the 25 random instances. While it is difficult to draw definitive claims for the small values of n accessible, we see an interesting split in the results. For low-rank ($n_A = 1, R = 2$) approximations, the number of iterations required seems to scale more favorably for our method than for both forms of tomography for both the SP and CCPS ansätze. The opposite appears to be true for the full- ϵ -rank approximation where it appears that even full tomography is a better strategy at $n = 3$ for both ansätze. This could plausibly be explained by the barren plateau phenomenon for learning random states that was proven in [22]; hence reconfirming that variational methods are not well suited to fully learning typical random states. On the other hand, with no shot noise and such small n , this could also be due to the presence of many local minima [71, 72]. Regardless of the cause, our simulations suggest that while our algorithm can learn small unstructured random states, it cannot scale beyond modest n .

3.3. XY model results

This discussion naturally raises the question of what happens when structure is present. This leads us to the study of thermal states in the Heisenberg XY model, given by

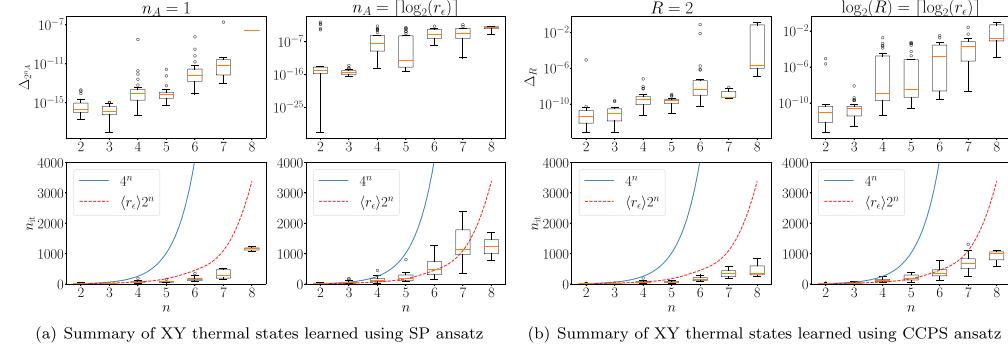


Figure 4. Compiling XY thermal states with idealized classical numerics. We summarize the XY thermal state results for the SP ansatz (left) and the CCPS ansatz (right). Each box plot bins the quartiles from 50 thermal states of the XY model where $J_i, K_i \sim \mathcal{N}(0, 1)$ are i.i.d. standard Gaussian random variables: 25 at a low temperature $\beta = 20$ and 25 at a higher temperature $\beta = 2$. Specifically, the orange line is the median, the box contains 75% of the runs, the top and bottom lines show the max and min, and circles represent outliers. The left column of each four-panel figure corresponds to a low-rank approximation whereas the right corresponds to a full ϵ -rank approximation. The first row shows how close the optimized state is to the best possible state, Δ_R , as a function of the number of qubits, n . The bottom row shows the number of iterations, n_{it} , it takes to perform the optimization alongside the naive scaling of full quantum state tomography 4^n and enhanced quantum sensing tomography $\langle r_e \rangle 2^n$ where the average is across the 25 random instances.

$$\rho_n^{(\text{XY})} := \frac{e^{-\beta H_{\text{XY}}}}{\text{Tr}[e^{-\beta H_{\text{XY}}}]}, \quad (33a)$$

$$H_{\text{XY}} := \sum_{i=1}^{n-1} J_i X_i X_{i+1} + K_i Y_i Y_{i+1}, \quad (33b)$$

where $J_i, K_i \sim \mathcal{N}(0, 1)$ are i.i.d. standard Gaussian random variables and $\beta = 1/k_B T$ is the inverse temperature. By controlling the temperature, we can control the ϵ -rank of the generated mixed states (see appendix D.2 for more details). Structurally, this model is clearly invariant under any global rotation of all spins, and the number of spins is constant. Thus, even with random coefficients, it exhibits important symmetries which allow us to greatly simplify our ansatz (see appendix B). For this reason, we are able to test from $n = 2$ to $n = 8$ qubits relatively easily, which is sufficient for an initial study of empirical resource scaling.

We tested our algorithm on 50 random XY chain thermal states (see equation (33a)) for each n : 25 at a large inverse temperature of $\beta = 20$ (i.e. low temperature regime) and 25 at a relatively smaller inverse temperature of $\beta = 2$ (i.e. moderate to high temperature regime). The results are shown in figure 4, whose format is the same as that in figure 3. Similar to the Bures results, we generated the results using noiseless classical optimization of an alternating layer ansatz with the Powell optimizer. Unlike Bures, however, the ansatz consisted of so-called Givens gates (see appendix B for exact definition) which respect the symmetries of the XY model.

Our first observation is that the performance (i.e. both Δ_R and n_{it}) does not have a noticeable dependence on β . For this reason, we chose to simplify the presentation of the results and combine all 50 states (for each n) into a single box plot. Note that this was not expected *a priori*. As discussed in appendix D.2, $\beta = 20$ corresponds to the limit when $\rho_n^{(\text{XY})}$ is approximately in the ground state of H_{XY} with a low ϵ -rank, whereas $\beta = 2$ samples intermediate to large ϵ -ranks. Hence, what we have found is that our optimization is insensitive to the underlying ϵ -rank of the target state in this case.

For the SP ansatz, the performance (as measured by Δ_R) is quite good, with a worst case of $n = 8, n_A = [\log_2 r_e], \Delta_R \approx 10^{-6}$. The same point for the CCPS ansatz only reached $\Delta_R \approx 10^{-4}$ in median performance, and indeed, the CCPS performance is noticeably worse across the entire data set. That is, when solving the same problem, the final Δ_R for the CCPS ansatz is often noticeably larger than the SP ansatz. However, even this worst median $\Delta_R \approx 10^{-4}$ is an acceptable ‘4 nines’ result (i.e. $C(\alpha^*, R)$ differs from $C(\alpha_{\text{opt}}, R)$ only in the fourth decimal place).

Across the entire data set, whether low or full rank and SP or CCPS, the actual number of iterations scales slightly better than the compressed tomography scaling of $\langle r_e \rangle 2^n$ and is substantially better than naïve tomography. For example, at $n = 8$, both low and high rank results use ≈ 56 times fewer iterations (relative to $4^n = 4^8$ for full naive tomography) with SP. However, this positive result must be considered along with the result that the quality of the solution, Δ_R , deteriorates with n . A fair summary of the result can be understood by setting an acceptable cut-off, δ_c . Supposing that $\delta_c = 10^{-3}$ (a ‘three nines criterion’), what our

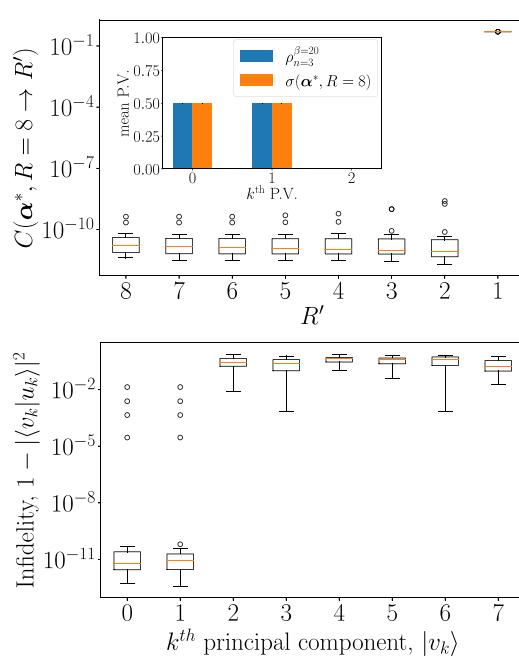


Figure 5. PCA and truncation with the CCPS ansatz. We learn a full rank ($R = 8$) approximation of 25 random three-qubit XY thermal states, achieving a median $\Delta_{R=8} \approx 10^{-10}$. Top: In the *main plot*, we show the Hilbert–Schmidt cost between the target and learned CCPS ansatz when truncating to keep the first R' states, i.e. $|\psi_k\rangle$ from $k \in \{1, \dots, R'\}$. In the *inset*, we show the distribution of principal values (i.e. ordered eigenvalues of ρ) for these XY thermal states alongside the found principal values of σ . Bottom: We show the infidelity between the k th principal component of ρ , $|\psi_k\rangle$, and the CCPS learned estimate $|u_k\rangle$.

data shows is that we can reach $\Delta_R < \delta_c$ in fewer iterations than both naïve tomography and compressed sensing tomography.

3.4. PCA and state compression example

As discussed, the CCPS ansatz (see equation (18)) takes the form of a convex combination of the R estimated principal components of ρ weighted by the associated principal values. Hence, learning $\sigma_{\text{CCPS}}(\alpha^*, R)$ is tantamount to performing PCA. As a consequence, we get the form of any $R' < R$ approximation from the rank R solution for free by truncation (see equation (31a)). We explore the practical meaning of these two statements by example in figure 5. In this case, we learn a full rank CCPS ansatz approximation of a random three qubit XY thermal state at low temperature, $\rho_{n=3}^{\beta=20}$. As before, we consider 25 such optimizations over randomly drawn coefficients (see equation (33a)).

In the top plot of figure 5, we explore the effect of truncation from $R = 8$ to $R' < R$ on the cost as a function of R' . At $R = 8$, the median cost is $C(\alpha^*, R) = \Delta_{R=8} \approx 10^{-10}$ which means we have successfully learned $\rho_{n=3}^{\beta=20}$. But even as we truncate down to $R' = 2$, the cost hardly changes—only jumping to a large value of 0.5 at $R' = 1$. This suggests that only two principal components are necessary to approximate $\rho_{n=3}^{\beta=20}$, and we corroborate this intuition by showing the principal values (or spectrum) of both $\rho_{n=3}^{\beta=20}$ and $\sigma(\alpha^*, R = 8)$ in the inset. Here, the bar-plot shows the median k th principal value as a function of k . For $\rho_{n=3}^{\beta=20}$, the values are $\lambda_0 \approx 1/2$, $\lambda_1 \approx 1/2$, and $\lambda_3 \approx 10^{-13}$, so $r_\epsilon = 2$ for $\epsilon > 10^{-13}$ for at least half of the instances (and all for $\epsilon > 10^{-4}$).

In the bottom plot of figure 5, we show the pure state infidelity between the k th principal component of $\rho_{n=3}^{\beta=20}$, $|\psi_k\rangle$, and the associated estimate contained in the CCPS ansatz, $|u_k\rangle = U_{\theta^*} |k\rangle$ as a function of k . Note that the k labels are ordered by decreasing principal value, i.e. $k = 0$ corresponds to the principal component with the largest principal value and so on. Clearly, the infidelity for $k = 0$ and $k = 1$ is very small $\sim 10^{-11}$ whereas the infidelity for $k \geq 2$ can be rather large. This again is due to $\rho_{n=3}^{\beta=20}$ having an effective ϵ -rank of two. But by plotting the pure state infidelity we have also made the notion of ‘learning the principal components of ρ ’ more explicit. Namely, a good approximation of ρ with the CCPS ansatz relies on having a high quality and explicit estimate of its important principal components as weighted by the relative importance of the principal values.

Finally, this discussion suggests an alternative way to use our algorithm as a means to find the approximate rank of an unknown state. By training for different values of R until $C \ll 1$ or C converges, we can estimate that $\text{rank}(\rho) \approx R$. This procedure also clarifies what we mean by claiming that our algorithm allows us to ‘learn a lower rank approximation that allows for more efficient processing.’ In this example, we

mean that learning a rank $R = 2$ approximation is sufficient, and therefore $\sigma_{CCPS}(\alpha^*, R = 2)$ is a low-rank approximation/compression of $\rho_{n=3}^{\beta=20}$.

4. Quantum hardware implementation

Finally, we consider the most important task for our algorithm: compiling a quantum state *on a quantum device*. In figure 6, we demonstrate the ability to learn two classes of quantum states on the IBM superconducting qubit devices with both ansatz choices. In figure 6(a), we successfully compile a random single qubit mixed state with both the SP and CCPS ansatz. Here, the random state is generated by tracing over one qubit in a two-qubit Haar random state

$$\rho_{HS} = \text{Tr}_A[U_{\text{Haar}}|00\rangle\langle 00|_{AB}U_{\text{Haar}}^\dagger]. \quad (34)$$

The HS subscript signifies that such a state is uniformly sampled from the HS metric [70]. Finding a compilation for a HS random state has a similar proof-of-principle goal to the Bures random states, but this class is easier to prepare on NISQ devices. We remark that the state ρ_{HS} is known *a priori* in this example in the sense that we supply the quantum device with U_{Haar} .

In figure 6(b), we learn *unknown* states generated by noisy state preparation which we call ‘NISQ’ states. In particular, we learn a noisy $|+\rangle$ state and a noisy $|\Phi^+\rangle$ state,

$$\tilde{\rho}_+ = \mathcal{E}_1(|+\rangle\langle +|) \quad (35a)$$

$$\tilde{\rho}_{\Phi^+} = \mathcal{E}_2(|\Phi^+\rangle\langle \Phi^+|), \quad (35b)$$

a single- and two-qubit state, respectively. For $\tilde{\rho}_+$, we apply a single Hadamard and then idle for the duration of 20 Hadamards. For $\tilde{\rho}_{\Phi^+}$, we follow the same procedure but then apply a final CNOT at the end. Under perfect conditions, these would generate the states $|+\rangle$ and $|\Phi^+\rangle$. Due to intrinsic ZZ cross-talk on superconducting devices [73–76] along with other secondary sources of noise, the qubits undergo a complicated dephasing process which we summarize as some unknown quantum channels \mathcal{E}_1 and \mathcal{E}_2 , respectively. For some sense of the strength of noise on these devices, we remark that $\text{Tr}[\tilde{\rho}_+^2] = 0.90(1)$ and $\text{Tr}[\tilde{\rho}_{\Phi^+}^2] = 0.78(4)$ with 1σ confidence intervals generated from ten bootstrapped tomography experiments. Since we do not know the quantum channels in advance, compiling the NISQ states is also tantamount to learning the states. For example, we can learn a low-depth circuit to prepare $\tilde{\rho}_+$, which is short enough to be unaffected by ZZ cross-talk. This serves as a permanent snapshot to probe $\tilde{\rho}_+$ even when \mathcal{E}_1 inevitably drifts due to two-level system and calibration effects.

With the states now defined, we discuss the results presented in figure 6 more closely. For each optimization we plot two costs, C_{shot} (dotted line) and $C_{\text{noiseless}}$ (solid line). The shot cost is computed on quantum hardware using 10^5 shots, and it is the cost we optimize over using the Powell optimizer. The noiseless cost is computed classically in post-processing for verification. For the random states this amounts to classically storing the circuits to prepare ρ and σ throughout the optimization and computing the cost with matrix operations. For the NISQ states, we rely on full quantum state tomography to compute ρ as a ‘trusted third party’ method since the states are generated by unknown noise. We terminate the optimization when either C_{shot} flattens for at least 10 iterations or 100 iterations are reached. In all cases, the final noiseless cost reaches $10^{-3} \leq C_{\text{noiseless}} \leq 10^{-2}$. We remark that this is consistent with the use of $N = 10^5$ shots since we expect a reported precision to scale as $\sim 1/\sqrt{N}$. In all cases, we show the result of learning a full rank approximation of the target state, so $\Delta_R = C_{\text{noiseless}}$ and hence $10^{-3} \leq \Delta_R \leq 10^{-2}$. For all but the $\tilde{\rho}_{\Phi^+}$ optimization we note that the noiseless cost is an order of magnitude lower than the cost evaluated on the hardware which is indicative of optimal parameter noise resilience [12]. That is, it suggests that the position of the global cost minimum of the cost landscape is (approximately) invariant under the action of noise.

Overall, our results show that we can successfully learn full-rank approximations of hardware relevant states. As we might expect from the Bures and XY state results, we can also learn lower rank approximations which we show explicitly in appendix F. The net result is very similar: we can learn lower rank approximations to within a precision of 10^{-3} to 10^{-2} , and as in the idealized classical experiments, it takes fewer iterations to learn lower rank approximations. Alternatively, we may choose to learn a high rank approximation of ρ with the CCPS ansatz and obtain lower rank approximations via truncation (as in section 3.4) which we also explore for the hardware data in appendix F.

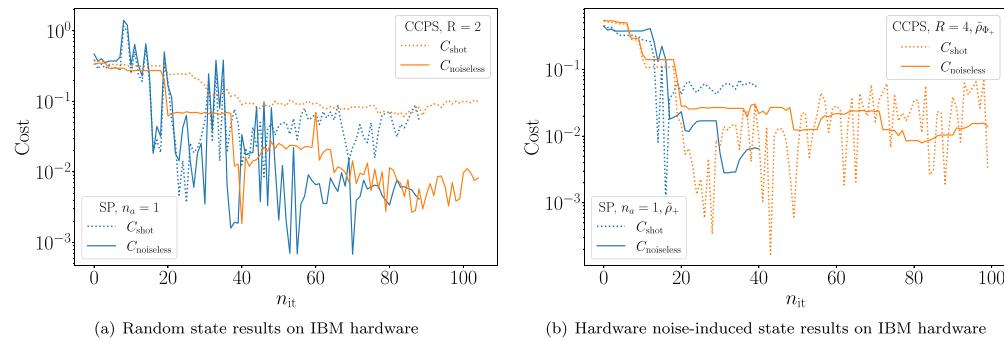


Figure 6. Hardware Implementation of QMSC We demonstrate the ability to compile quantum states on IBM superconducting devices by evaluating the cost function on them. In (a), we compile random single qubit states generated by equation (34) with both the SP and CCPS ansatz successfully. In (b), we learn two hardware-noise induced states which are described in equations (35a) and (35b) using the SP and CCPS ansatz, respectively. Optimizations were performed on the seven-qubit devices `ibmq_casablanca` (random SP result) and `ibmq_jakarta` (all others) using $N = 10^5$ shots to evaluate C_{shot} (dotted line). A noiseless cost $C_{noisless}$ was also computed classically for verification. In each case, $10^{-3} \leq C_{noisless} \leq 10^{-2}$ is reached which denotes a successful compiling given 10^5 shots.

5. Discussion and outlook

We have presented an algorithm to learn an unknown mixed state ρ . In particular, we have developed a procedure to learn a rank- R approximation to ρ , where $\text{rank}(\rho) \geq R$ is assumed but not essential. Put precisely, our algorithm is a practical variational way to solve the QLRAP [18]. Applications of this algorithm are numerous and include PCA, state compression, learning noise-induced states, and uploading of states onto quantum computers, as described in figure 1.

We considered two ansatz constructions. If the purification ansatz is chosen, the end result is a unitary $V(\theta^*)|0\rangle^{\otimes n} \otimes |0\rangle^{\otimes n_A}$ generates a purification of the rank- R approximation of ρ . By tracing out the n_A ancilla qubits, we get the desired rank R mixed state. For the CCPSs ansatz, the final output is a classical vector \mathbf{p} containing the R principal values of ρ and a unitary $U(\theta^*)$ such that $U(\theta^*)|i\rangle$ for $i \in \{1, \dots, R\}$ gives the R principal vectors of ρ .

Our numerical simulations and hardware implementations indicate that our algorithm works well for a variety of state ensembles, including random XY spin chain thermal states at arbitrary temperatures and unknown states generated by hardware noise on superconducting qubit devices. Unsurprisingly, we found that learning XY thermal states was easier than random states because the additional structure of the problem opened up the possibility of using simpler ansätze. Additionally, while the SP ansatz performs better in numerical simulations (because its optimization problem is simpler), the CCPS ansatz, as expected, allows for larger hardware implementations because it requires fewer qubits.

For both ansätze, our algorithm provides a means of compressing the target state when $\text{rank}(\rho) > R$. For the purification ansatz, the reduction is in terms of the number of qubits; for the convex combination ansatz, the reduction is in terms of the number of pure states required to simulate the effect of the state. While the compression of states [16, 77–81], and indeed data sets encoded in states [82], has been explored previously, much of the compression-based literature focuses on finding the compressed state via maximizing the degree to which the original state can be reconstructed via a successful decompression process [83].

A particularly timely application of our algorithm is for quantum PCA. While quantum PCA was originally proposed to have an exponential speedup over classical methods for low rank states [35], it was later dequantized for the case of classical data analysis [84], reducing speedups for this case to being modest ones [85]. However, recently it was shown that these dequantization arguments break down for quantum data analysis [86] and that quantum PCA for quantum data can indeed achieve an exponential quantum speedup [87]. Moreover a simple method to encode the covariance matrix into a density matrix was recently proposed [88], making quantum PCA an easily accessible application for near-term quantum computers. Hence, our approach for extracting the principal components of a density matrix is especially timely, in the quest for near-term quantum advantage.

In this article we have focused on coherent access models for quantum compilation. That is, computing the cost using the Loschmidt echo or SWAP test requires coherent interaction between the state we wish to compile and the device on which we wish to compile it. For this to be possible, the target quantum state either needs to be already prepared (potentially by some unknown process) on the quantum computer or we require a quantum sensor to mediate the interaction between the target quantum system and the quantum computer. The former is practically viable and reasonable to assume for applications such as learning

noise-induced states or state compression. However, further developments in quantum sensors will be required to upload the unknown quantum state of an experimental system to a quantum computer in this coherent access model.

An alternative approach would be to use Clifford shadow tomography [6–8] to efficiently compute the overlap between the target state and a large set of guess compilations via independent measurements and then classical post-processing. This incoherent version of quantum compilation could be applied in situations where the unknown state is prepared on a platform that is very different from the hardware on which we wish to compile it. In this manner, this approach opens up new techniques for uploading the quantum state of experimental systems to quantum hardware. In figure 1 we sketch the difference between the coherent and incoherent access models. Additionally, [6] provides a means of upper bounding the copy complexity of such a compilation task. In the case of efficiently preparable target states and ansatz states, i.e. those that can be prepared via a circuit with $T \in \mathcal{O}(\text{poly}(n))$ local gates, then $\mathcal{O}(T \log(T/\epsilon)/\epsilon^2) \leq \tilde{\mathcal{O}}(\text{poly}(n)/\epsilon^2)$ copies of the target state ρ suffices to approximately evaluate the cost and its partial derivatives to precision ϵ arbitrarily often [89]. Thus $\tilde{\mathcal{O}}(\text{poly}(n)/\epsilon^2)$ copies of the target ρ in theory suffice to compile it to precision $\tilde{\mathcal{O}}(\epsilon)$.

The framework investigated in this article for mixed state compilation may more generally be applied to the compiling of quantum channels. This follows from the fact that a channel may be represented via its Choi representation as a mixed state. That is, a channel may be fully characterized via the mixed state generated by applying a quantum channel to one half of a Bell state. Therefore one means of compiling a quantum channel would be to minimize the HS distance between the Choi state corresponding to a target channel and an ansatz mixed state, formed by applying a parameterized channel to half a Bell state. In this sense, our algorithm further open up new avenues for learning unknown quantum processes.

Data availability statement

The data that support the findings of this study are openly available at the following URL/DOI: <https://doi.org/10.5281/zenodo.7039676>.

The data and source code that support the findings of this study are openly available in a Zenodo repository [59], a static and citable version of an available Github repository.

Acknowledgments

N E was supported by the U.S. Department of Energy (DOE) Computational Science Graduate Fellowship under Award Number DE-SC0020347. E M B acknowledges support from The Engineering and Physical Sciences Research Council (EPSRC) in the UK. A U S acknowledges support from the Center for Computation and Technology (CCT) at Louisiana State University. M M W acknowledges support from the National Science Foundation under Grant No. 1907615. A T S acknowledges support from the QSC (see below). P J C was initially supported by the Los Alamos National Laboratory (LANL) ASC Beyond Moore's Law Project. P J C also acknowledges later funding from the Laboratory Directed Research and Development (LDRD) program of LANL under Project Number 20210116DR. Z H acknowledges initial support from the LANL Mark Kac Fellowship and subsequent support from the Sandoz Family Foundation-Monique de Meuron program for Academic Promotion. This work was also supported by the Quantum Science Center (QSC), a National Quantum Information Science Research Center of the U.S. Department of Energy (DOE).

Appendix A. Local cost function analysis (HS)

As discussed briefly in section 2, barren plateaus form a barrier for trainability of VQAs for large-scale problems. Here, we formulate and analyze different proposals of localized versions of the HS distance.

A.1. Marginal local cost function

A simple, naïve choice in 1-local cost could be formulated in terms of the HS distance between the local reduced states. We call this cost function the marginal 1-local cost function, and it takes the same form for both the SP and CCPSS ansatz.

Definition 1. Given two quantum states ρ and σ , let us define a *Marginal 1-local cost* as follows:

$$C_M^{(1)}(\rho, \sigma) := \frac{1}{n} \sum_{j=1}^n \|\rho_j - \sigma_j\|_2^2, \quad (\text{A1})$$

where $\rho_j = \text{Tr}_{\bar{j}}[\rho]$ and $\sigma_j = \text{Tr}_{\bar{j}}[\sigma]$, with \bar{j} denoting all system qubits except for qubit j .

This can be measured using the methods outlined in section 2. (Note, in contrast to the main text, in this appendix we stress the dependence of the cost on the target state ρ and guess state σ .)

A successful candidate for a local cost function needs to remain faithful to the global cost, i.e. the minimal value of the cost corresponds to the case when $\rho = \sigma$. This cost function is trivially faithful for learning product states using a product state ansatz, i.e. if $\sigma = \bigotimes_{j=1}^n \sigma_j$ and $\rho = \bigotimes_{j=1}^n \rho_j$. However, it does not take much thought to see that this local cost function is not faithful more generally. In simple terms, different entangled states may have the same reduced states and so the local cost may vanish even when the global states are non-identical. For example, consider the case in which ρ and σ are orthogonal Bell states, i.e. $\rho = |\psi_+\rangle\langle\psi_+|$ and $\sigma(\theta, n_A) = |\psi_-\rangle\langle\psi_-|$. Now, in this case, the local cost vanishes because the reduced states are all maximally mixed

$$C_M^{(1)}(\rho, \sigma) = \frac{1}{n} \sum_{j=1}^n \|\rho_j - \sigma_j\|_2^2 \quad (A2)$$

$$= \frac{1}{n} \sum_{j=1}^n \|\mathbb{1}/2 - \mathbb{1}/2\|_2^2 = 0. \quad (A3)$$

However, $|\psi_+\rangle\langle\psi_+| \neq |\psi_-\rangle\langle\psi_-|$ and so the local cost is not faithful. In contrast, the equivalent global cost, that is the HS distance between $|\psi_+\rangle\langle\psi_+|$ and $|\psi_-\rangle\langle\psi_-|$, is non zero,

$$D_{\text{HS}}(\rho, \sigma) := \sqrt{\text{Tr}[\rho^2] + \text{Tr}[\sigma^2] - 2\text{Tr}[\rho\sigma]} \quad (A4)$$

$$= \sqrt{1+1-0} = \sqrt{2}, \quad (A5)$$

as it should be.

The lack of faithfulness does not necessarily preclude us using it for training. We could start training on the 1-local cost and as we approach the solution, add in k -local cost terms where we look at the distance between the reduced states on k qubits. That is, we propose to use a cost of the following form.

Definition 2. Given two quantum states ρ and σ , and some (possibly multiple) partitioning(s)¹² \mathcal{P}_k of the set of n qubits into N_k subsets of at most $k < n$ qubits, where at least one set has cardinality k , let us define the *Generalized Marginal Cost* as follows:

$$C_M^{\text{gen}}(\rho, \sigma) := \sum_{k=1}^n \alpha_k(t) C_M^{(k, \mathcal{P}_k)}(\rho, \sigma), \quad (A6)$$

where

$$C_M^{(k, \mathcal{P}_k)}(\rho, \sigma) := \frac{1}{N_k} \sum_{i=1}^{N_k} \left\| \rho_i^{(k)} - \sigma_i^{(k)} \right\|_2^2, \quad (A7)$$

with $\{\rho_i^{(k)}, \sigma_i^{(k)}\}_{i=1}^{N_k}$ the set of marginals of ρ and σ as determined by \mathcal{P}_k , and the iteration dependent weightings $\alpha_k(t) \geq 0$ satisfying $\sum_k \alpha_k(t) = 1$.

We note that an analysis of the general behavior of $C_M^{(k)}(\rho, \sigma)$ is non-trivial as the behavior depends on the partitions \mathcal{P}_k . One (simple) potential choice of partition is to only allow k such that $k \mid n$, and choose N_k sets of k qubits. In this case, for states that are k -local product states on the chosen subsets, $C_M^{(k)}(\rho, \sigma)$ would be trivially faithful. By starting with the 1-local cost, i.e. with $\alpha_1 = 1$, and by slowly ramping to the global cost $C_M^{(n)} \equiv C_M^{(n)}(\rho, \sigma)$ by increasing the weightings of the various α_k terms until $\alpha_n = 1$, it should be possible to train to the minimum of the global cost. We expect this proposal to prove most useful for learning mixed states with relatively localized entanglement.

¹² Choosing partitions is relatively arbitrary. One simple choice would be to partition the n qubits into $\lfloor n/k \rfloor$ subsets of k qubits and one subset of $n \% k$ qubits. It would be natural to assign the qubits to the subsets to minimize the distance between qubits in a subset as determined by the target state's/hardware's geometry. However, one can imagine sets of more complex partitions including those where each qubit belongs to multiple subsets. For example, you could consider the set of all possible subsets of k qubits; however, this would become resource intensive to measure in practice for $1 \ll k \ll n$.

A.2. Local measurements on global states

An alternative approach to formulating a local cost function is to continue to use the full n -qubit input states, but replace the global measurements with *local measurements* (similarly to [10]). As each ansatz includes specifically prescribed measurement operators, we treat each of our ansätze separately, and derive provable guarantees of faithfulness for special cases.

A.2.1. Local costs for CCPS ansatz

Consider the global HS cost function in the CCPS ansatz framework:

$$C_{\text{CCPS}}(\boldsymbol{\alpha}, R) = \text{Tr}[\rho^2] + \sum_i p_\phi(i)^2 - 2 \sum_i p_\phi(i) \text{Tr} \left[U_\theta^\dagger \rho U_\theta H_G^{(i)} \right], \quad (\text{A8})$$

where we write $H_G^{(i)} := |i\rangle\langle i|_S$ to emphasize the globality of the measurement. Since the purity of the target state is not optimized and the purity of the guess state is optimized fully classically, these terms can be left global without impeding trainability. However, for the overlap term, we seek to replace the global $H_G^{(i)}$ term with some set of local measurements. There is freedom in how we may choose our local measurements, but one simple approach is to choose the average of all 1-local measurements, as follows

Definition 3. For the CCPS ansatz, we can define a *Local cost function* as

$$C_L^{\text{CCPS}} = \text{Tr}[\rho^2] + \sum_i p_\phi(i)^2 - 2 \sum_i p_\phi(i) \text{Tr}_S \left[U_\theta^\dagger \rho U_\theta H_L^{\text{CCPS}(i)} \right], \quad (\text{A9})$$

with

$$H_L^{\text{CCPS}(i)} = \frac{1}{n} \sum_{j=1}^n |i_j\rangle\langle i_j|_{S_j} \otimes \mathbb{1}_{S_{\bar{j}}}, \quad (\text{A10})$$

where, as in (18), $|i\rangle = \bigotimes_{j=1}^n |i_j\rangle$ denotes an element of the computational basis of n qubits, and $|i_j\rangle$ denotes the j th bit of the bit string i .

In order to prove faithfulness in the case of pure states, we first state and re-derive a result introduced in [10].

Proposition 1. Let ρ be an arbitrary quantum state, and let U_θ be a parameterized unitary matrix. Denoting $\rho_U = U_\theta^\dagger \rho U_\theta$, and $H_L^{(i)} = H_L^{\text{CCPS}(i)}$, we have that

$$1 - \text{Tr} \left[\rho_U H_L^{(i)} \right] \leq 1 - \text{Tr} \left[\rho_U H_G^{(i)} \right] \leq n \left(1 - \text{Tr} \left[\rho_U H_L^{(i)} \right] \right). \quad (\text{A11})$$

Proof. We can write $H_L^{(i)} = \frac{1}{n} \sum_j^n H_{L,j}^{(i)}$, where

$$H_{L,j}^{(i)} = |i_j\rangle\langle i_j|_{S_j} \otimes \mathbb{1}_{S_{\bar{j}}} \quad (\text{A12})$$

are projectors that mutually commute. Note that $\prod_{j=1}^n H_{L,j}^{(i)} = H_G^{(i)}$. We can associate events E_j with the projectors $H_{L,j}$ such that $\text{Pr}[E_j] = \text{Tr} \left[\rho_U H_{L,j}^{(i)} \right]$. Then, $\text{Tr} \left[\rho \prod_{j=1}^n H_{L,j}^{(i)} \right] = \text{Pr} \cap_{j=1}^n E_j$. Recall, from basic probability theory, that for any set of events $\mathcal{A} := \{A_1, A_2, \dots, A_n\}$, it holds that

$$\text{Pr} \left[\bigcup_{i=1}^n A_i \right] \geq \frac{1}{n} \sum_{i=1}^n \text{Pr}[A_i]. \quad (\text{A13})$$

Choosing $A_i = \overline{E_j}$, we see

$$\begin{aligned} \text{Pr} \left[\bigcup_{j=1}^n \overline{E_j} \right] &\geq \frac{1}{n} \sum_{i=1}^n \text{Pr} \left[\overline{E_j} \right] \\ \implies 1 - \text{Pr} \left[\bigcap_{j=1}^n E_j \right] &\geq \frac{1}{n} \sum_{j=1}^n (1 - \text{Pr}[E_j]) \\ \implies 1 - \text{Pr} \left[\bigcap_{j=1}^n E_j \right] &\geq 1 - \frac{1}{n} \sum_{j=1}^n \text{Tr} \left[\rho_U H_{L,j}^{(i)} \right] \\ \implies 1 - \text{Tr} \left[\rho_U H_G^{(i)} \right] &\geq 1 - \text{Tr} \left[\rho_U H_L^{(i)} \right]. \end{aligned} \quad (\text{A14})$$

This is precisely the first desired inequality $1 - \text{Tr}[\rho_U H_{\text{L}}^{(i)}] \leq 1 - \text{Tr}[\rho_U H_{\text{G}}^{(i)}]$. To prove the remaining inequality, observe that, via the union bound, we have

$$\begin{aligned} \text{Pr}\left[\bigcup_{j=1}^n \overline{E_j}\right] &\leq \sum_{j=1}^n \text{Pr}[\overline{E_j}] \\ \implies 1 - \text{Pr}\left[\bigcap_{j=1}^n E_j\right] &\leq \sum_{j=1}^n (1 - \text{Pr}[E_j]) \\ \implies 1 - \text{Tr}[\rho_U H_{\text{G}}^{(i)}] &\leq n(1 - \text{Tr}[\rho_U H_{\text{L}}^{(i)}]). \end{aligned} \quad (\text{A15})$$

Thus, $1 - \text{Tr}_{\text{SB}}[\rho_U H_{\text{L}}^{(i)}] \leq 1 - \text{Tr}_{\text{SB}}[\rho_U H_{\text{G}}^{(i)}] \leq n(1 - \text{Tr}_{\text{SB}}[\rho_U H_{\text{L}}^{(i)}])$ as required. \square

Proposition 2. $C_{\text{L}}^{\text{CCPS}}$ is faithful for pure states, and ‘close to faithful’ for states with low impurity. Specifically we have that

$$nC_{\text{L}}^{\text{CCPS}} \geq C_{\text{G}} - (n-1)(\text{Impurity}(\rho) + \text{Impurity}(\sigma)), \quad (\text{A16})$$

where $\text{Impurity}(X) := 1 - \text{Tr}[X^2]$ for $X = \rho$ and $X = \sigma$ and we write $C_{\text{G}} = \|\rho - \sigma\|_2^2$ to emphasize the globality of the standard HS distance cost.

Remark 3. It follows from proposition 2, that if $C_{\text{L}}^{\text{D}} = 0$ then

$$(n-1)(\text{Impurity}(\rho) - \text{Impurity}(\sigma)) \geq \|\rho - \sigma\|_2^2. \quad (\text{A17})$$

That is, if the target and trained state are pure, i.e. $\text{Impurity}(\sigma) = \text{Impurity}(\rho) = 0$, we have that $C_{\text{L}} = 0$ entails that $C_{\text{G}} = 0$. More generally, if the impurities of the target and trained states are low, $C_{\text{L}} = 0$ entails that C_{G} is small.

Proof. From (A11), we have

$$\begin{aligned} 1 - \text{Tr}[\rho_U H_{\text{G}}^{(i)}] &\leq n(1 - \text{Tr}[\rho_U H_{\text{L}}^{(i)}]) \\ \implies -\text{Tr}[\rho_U H_{\text{G}}^{(i)}] &\leq (n-1) - n\text{Tr}[\rho_U H_{\text{L}}^{(i)}] \\ \implies -2 \sum_i p_{\phi}(i) \text{Tr}[\rho_U H_{\text{G}}^{(i)}] &\leq -2n \sum_i p_{\phi}(i) \text{Tr}[\rho_U H_{\text{L}}^{(i)}] + 2(n-1) \sum_i p_{\phi}(i), \end{aligned} \quad (\text{A18})$$

where $2(n-1) \sum_i p_{\phi}(i) = 2(n-1)$. Adding the purity terms to both sides gives

$$\begin{aligned} C_{\text{G}} &\leq \text{Tr}[\rho^2] + \text{Tr}[\sigma^2] - 2n \sum_i p_{\phi}(i) \text{Tr}[\rho_U H_{\text{L}}^{(i)}] + 2(n-1) \\ \implies C_{\text{G}} &\leq nC_{\text{L}}^{\text{CCPS}} + (n-1)(2 - \text{Tr}[\rho^2] - \text{Tr}[\sigma^2]), \end{aligned} \quad (\text{A19})$$

which can be rewritten as

$$nC_{\text{L}}^{\text{CCPS}} \geq C_{\text{G}} - (n-1)(\text{Impurity}(\rho) + \text{Impurity}(\sigma)). \quad (\text{A20})$$

Accordingly, if $C_{\text{L}}^{\text{CCPS}} = 0$, then

$$C_{\text{G}} \leq (n-1)(\text{Impurity}(\rho) + \text{Impurity}(\sigma)). \quad (\text{A21})$$

Therefore if the purity of the target and trained states are zero, the cost is faithful; i.e. $C_{\text{L}}^{\text{CCPS}} = 0$ implies $C_{\text{G}} = 0$. Similarly, for high purity states the cost is approximately faithful, i.e. $C_{\text{L}}^{\text{CCPS}} = 0$ implies C_{G} is small. \square

Similar to the generalized marginal cost defined earlier, we can construct an extension of the low impurity local cost function, such that it is local on k qubits, for $k \leq n$.

One simple approach is to perform n/k measurements, on k qubits at a time. This naturally restricts us to only choosing k such that $k \mid n$. Defining

$$H_{\text{L}}^{k(i)} := \frac{1}{(n/k)} \sum_{m=1}^{n/k} H_{\text{L}_m}^{k(i)}, \quad (\text{A22})$$

with

$$H_{L_m}^{k(i)} := |i_{\mathcal{P}_m}\rangle\langle i_{\mathcal{P}_m}|_{\mathcal{P}_m} \otimes \mathbb{1}_{\overline{\mathcal{P}_m}}, \quad (\text{A23})$$

where \mathcal{P}_m contains the k indices of the k qubits being measured over by the m th operator, such that $\mathcal{P}_1 \cup \dots \cup \mathcal{P}_{n/k}$ spans $\{1, \dots, n\}$; we can define a k -local cost function as follows:

Definition 4. For $k \mid n$, we have the k -local cost function

$$C_L^k = \text{Tr}[\rho^2] + \sum_i p_\phi(i)^2 - 2 \sum_i p_\phi(i) \text{Tr}_S \left[U_\theta^\dagger \rho U_\theta^\dagger H_{L_m}^{k(i)} \right]. \quad (\text{A24})$$

Proposition 4. Let ρ be an arbitrary quantum state, and let U_θ be a parameterized unitary matrix. Denoting $\rho_U = U_\theta^\dagger \rho U_\theta$, we have that

$$1 - \text{Tr} \left[\rho_U H_L^{k(i)} \right] \leq 1 - \text{Tr} \left[\rho_U H_G^{(i)} \right] \leq \frac{n}{k} \left(1 - \text{Tr} \left[\rho_U H_L^{k(i)} \right] \right). \quad (\text{A25})$$

Proof. Similarly to Proposition 1, we can associate events E_i with the projections $H_{L_m}^{k(i)}$ such that $\text{Pr}[E_i] = \text{Tr} \left[\rho_U H_{L_m}^{k(i)} \right]$. Then, $\text{Tr} \left[\rho_U \prod_{m=1}^{n/k} H_{L_m}^{k(i)} \right] = \text{Pr} \left[\bigcup_{m=1}^{n/k} E_m \right]$. We have that

$$\begin{aligned} \text{Pr} \left[\bigcup_{m=1}^{n/k} \overline{E_m} \right] &\geq \frac{1}{(n/k)} \sum_{m=1}^{n/k} \text{Pr} \left[\overline{E_m} \right] \\ \implies 1 - \text{Pr} \left[\bigcap_{m=1}^{n/k} E_m \right] &\geq \frac{1}{(n/k)} \sum_{m=1}^{n/k} (1 - \text{Pr}[E_m]) \\ \implies 1 - \text{Tr} \left[\rho_U H_G^{(i)} \right] &\geq 1 - \frac{1}{(n/k)} \sum_{m=1}^{n/k} \text{Tr} \left[\rho_U H_{L_m}^{k(i)} \right], \end{aligned} \quad (\text{A26})$$

forming one side of our inequality. We also have

$$\begin{aligned} \text{Pr} \left[\bigcup_{m=1}^{n/k} \overline{E_m} \right] &\leq \sum_{m=1}^{n/k} \text{Pr} \left[\overline{E_m} \right] \\ \implies 1 - \text{Pr} \left[\bigcap_{m=1}^{n/k} E_m \right] &\leq \sum_{m=1}^{n/k} (1 - \text{Pr}[E_m]) \\ \implies 1 - \text{Tr} \left[\rho_U H_G^{(i)} \right] &\leq \frac{n}{k} \left(1 - \text{Tr} \left[\rho_U H_L^{k(i)} \right] \right). \end{aligned} \quad (\text{A27})$$

Thus,

$$1 - \text{Tr} \left[\rho_U H_L^{k(i)} \right] \leq 1 - \text{Tr} \left[\rho_U H_G^{(i)} \right] \leq \frac{n}{k} \left(1 - \text{Tr} \left[\rho_U H_L^{k(i)} \right] \right). \quad (\text{A28})$$

□

Proposition 5. The k -local cost function C_L^k is faithful for pure states, and ‘close to faithful’ for states with low purity, with this closeness increasing with k . Specifically we have that

$$\frac{n}{k} C_L^k \geq C_G - \left(\frac{n}{k} - 1 \right) (\text{Impurity}(\rho) + \text{Impurity}(\sigma)). \quad (\text{A29})$$

Remark 6. It follows from proposition 5, that if $C_L^k = 0$, then

$$\left(\frac{n}{k} - 1 \right) (\text{Impurity}(\rho) + \text{Impurity}(\sigma)) \geq \|\rho - \sigma\|_2^2. \quad (\text{A30})$$

Proof. The proof is entirely analogous to that for proposition 2 but with $n \rightarrow n/k$. □

Comparing this to the inequality found for the 1-local cost (A16), we find that the k -local cost is ‘closer to faithful’ at low impurities than the 1-local cost. It becomes increasingly faithful as k tends to n and, trivially, perfectly faithful for $k = n$. Thus similarly to the marginal local cost, we could start training on the 1-local cost and as we approach the solution, add in k -local cost terms to drive the ansatz towards the global minimum.

A.2.2. Local costs for SP ansatz

Consider our global HS cost function, in the SP ansatz framework:

$$C_{\text{SP}}(\rho, \sigma(\boldsymbol{\theta}, n_A)) = \text{Tr}[\rho^2] + \text{Tr}[\sigma(\boldsymbol{\theta}, n_A)^2] - 2\text{Tr}[\rho\sigma(\boldsymbol{\theta}, n_A)]. \quad (\text{A31})$$

Without loss of generality, we can also express ρ via its purification for analysis, i.e.

$$\rho = \text{Tr}_A \left[V_\rho (|0\rangle\langle 0|)^{\otimes(n+n_A)} V_\rho^\dagger \right], \quad (\text{A32})$$

where V_ρ is the purifying unitary associated with the target state ρ . Recall

$$\text{Tr}[\rho\sigma] = d_A \text{Tr}_{SA} \left[U_{\boldsymbol{\theta}}^\dagger \left(\rho \otimes \frac{\mathbb{1}_A}{d_A} \right) U_{\boldsymbol{\theta}} (|0\rangle\langle 0|)^{\otimes(n+n_A)} \right] \quad (\text{A33})$$

where d_A is the dimension of the environment system.

In order to construct a local cost, we can replace the global $(|0\rangle\langle 0|)^{\otimes(n+n_A)}$ projector in each term with a local measurement. The structure of this local measurement may, again, be chosen freely. In the case of learning ρ completely, (i.e. $R = \text{rank}(\rho)$), one simple approach is to replace the global measurement with the average of all measurements local to one system qubit and one environment qubit, i.e.

$$H_L^D = \frac{1}{n} \sum_{j=1}^n |0\rangle\langle 0|_{S_j} \otimes \mathbb{1}_{S_{\bar{j}}} \otimes |0\rangle\langle 0|_{A_j} \otimes \mathbb{1}_{A_{\bar{j}}}, \quad (\text{A34})$$

with \bar{j} denoting all qubits other than qubit j . We dub this the *doubly-local SP Hamiltonian*.

When seeking to learn a low-rank approximation (i.e. compression) of ρ , we need not make the measurement local on the ancilla. In this case we can use the following *singly-local SP Hamiltonian*,

$$H_L^S = \frac{1}{n} \sum_{j=1}^n |0\rangle\langle 0|_{S_j} \otimes \mathbb{1}_{S_{\bar{j}}} \otimes |0\rangle\langle 0|_A. \quad (\text{A35})$$

Definition 5. For the SP ansatz, we can define the *Doubly- and Singly-local costs*,

$$C_L^X = c_L^X(V_\rho, \rho) + c_L^X(U_{\boldsymbol{\theta}}, \sigma) - 2c_L^X(U_{\boldsymbol{\theta}}, \rho), \quad (\text{A36})$$

where

$$c_L^X(U_{\boldsymbol{\theta}}, \rho) := d_A \text{Tr}_{SA} \left[U_{\boldsymbol{\theta}}^\dagger (\rho_S \otimes \mathbb{1}/d_A) U_{\boldsymbol{\theta}} H_L^X \right], \quad (\text{A37})$$

with the freedom to choose arbitrary H_L^X , e.g. H_L^D or H_L^S as defined above.

The terms $c_L^X(U_{\boldsymbol{\theta}}, \sigma)$ and $c_L^X(U_{\boldsymbol{\theta}}, \rho)$ can be measured using Loschmidt-echo type circuits (as discussed in section 2) but with the global measurements replaced with local ones. However, as in general one will not have access to V_ρ , it is generally not possible to measure $c_L^X(V_\rho, \rho)$. Nonetheless, as this term remains constant throughout and does not contribute to the gradient (discussed further in appendix C) it can be neglected without effecting the optimization procedure.

The methods used to prove faithfulness of C_L^{CCPS} for pure states do not carry over for the singly- and doubly-local SP costs, due to the factor of d_A . Thus, faithfulness for pure states for C_L^D and C_L^S remains an open question. However, we can prove faithfulness for tensor-product states in the SP picture.

Proposition 7. C_L^D is faithful for tensor-product states.

Proof. In the case of tensor-product mixed states, we can write $\rho = \bigotimes_{j=1}^n \rho_{S_j}$, $U_{\theta} = \bigotimes_{j=1}^n U_{S_j A_j}^{\theta}$. Thus the overlap term can be written as

$$\begin{aligned}
 C_L^D(U_{\theta}, \rho) &= d_A \text{Tr}_{SA} \left[\left(\bigotimes_{k=1}^n U_{S_k A_k}^{\theta \dagger} \left(\bigotimes_{k=1}^n \rho_{S_k} \otimes \mathbb{1}_A / d_A \right) \bigotimes_{k=1}^n U_{S_k A_k}^{\theta} \right) H_L^D \right] \\
 &= \frac{d_A}{n} \sum_{j=1}^n \text{Tr}_{SA} \left[\left(\bigotimes_{k=1}^n U_{S_k A_k}^{\theta \dagger} \left(\bigotimes_{k=1}^n \rho_{S_k} \otimes \mathbb{1}_A / d_A \right) \bigotimes_{k=1}^n U_{S_k A_k}^{\theta} \right) (|0\rangle\langle 0|_{S_j} \otimes \mathbb{1}_{S_j} \otimes |0\rangle\langle 0|_{A_j} \otimes \mathbb{1}_{A_j}) \right] \\
 &= \frac{1}{n} \sum_{j=1}^n \text{Tr}_{S_j A_j} \left[\left(U_{S_j A_j}^{\theta \dagger} (\rho_{S_j} \otimes \mathbb{1}_{A_j}) U_{S_j A_j}^{\theta} \right) (|0\rangle\langle 0|_{S_j A_j}) \right] \text{Tr}_{S_j A_j} \left[\left(\bigotimes_{k \neq j} \rho_{S_k} \otimes \mathbb{1}_{A_j} \right) (\mathbb{1}_{S_j A_j}) \right] \\
 &= \frac{d_A}{2n} \sum_{j=1}^n \text{Tr}_{S_j A_j} \left[U_{S_j A_j}^{\theta \dagger} (\rho_{S_j} \otimes \mathbb{1}_{A_j}) U_{S_j A_j}^{\theta} |0\rangle\langle 0|_{S_j A_j} \right] \\
 &= \frac{d_A}{2n} \sum_{j=1}^n \text{Tr}[\rho_{S_j} \sigma_{S_j}].
 \end{aligned} \tag{A38}$$

Thus, in this case, we have that

$$\begin{aligned}
 C_L^D &= \frac{d_A}{2n} \left(\sum_{j=1}^n \text{Tr}[\rho_j \rho_j] + \sum_{j=1}^n \text{Tr}[\sigma_j \sigma_j] - 2 \sum_{j=1}^n \text{Tr}[\rho_j \sigma_j] \right) \\
 &\propto \sum_{j=1}^n \|\rho_j - \sigma_j\|_2^2,
 \end{aligned} \tag{A39}$$

which vanishes if and only if $\rho_j = \sigma_j$ for all j , i.e. assuming ρ and σ are tensor-product states, if $\rho = \sigma$. \square

In this subsection, we have shown how by introducing local measurements directly into our cost function, we are able to construct local cost functions. For the case of the CCPS ansatz this construction is provably faithful for pure states and approximately faithful for high purity states. For the case of the SP ansatz the construction is faithful for product states. Thus we expect these costs to prove useful for learning mixed states with relatively low impurities and low entanglement respectively. However, for target states that are highly entangled and/or mixed, we are unable to provide guarantees on the behavior of the cost function. In these cases, the function no longer resembles a distance measure, as positivity cannot be guaranteed. Thus the construction of a truly faithful, yet entirely local equivalent to the HS distance remains an open question.

However, in practice one may create a cost that is both faithful and exhibits non vanishing gradients by taking a linear combination of the absolute value of the local cost and the global cost, i.e. by training on a cost of the form $\alpha|C_L| + (1 - \alpha)C_{\text{Global}}$ for some choice of the local cost C_L and $0 \leq \alpha \leq 1$. By tuning α such that it is (close to) one at the start of the optimization and (close to) zero at later stages of the optimization, it should be possible to steer towards the global minimum.

Appendix B. Circuit ansätze particulars

B.1. Summarizing the circuits in our algorithm

We summarize the circuits used to evaluate our cost function to clarify in detail how one can implement our algorithm. In particular, we provide three circuits which sample the three terms in our cost function,

$$C(\alpha, R) \equiv \text{Tr}[\rho^2] + \text{Tr}[\sigma(\alpha, R)^2] - 2\text{Tr}[\rho\sigma(\alpha, R)] \tag{B1}$$

for both the SP ansatz and the CCPSs ansatz.

B.1.1. SP ansatz circuits

The SP ansatz generates an n qubit mixed state by applying a unitary on $n + n_A$ qubits and tracing out the n_A ancilla,

$$\sigma_{\text{SP}}(\theta, n_A) := \text{Tr}_A[U_{\theta}(|0\rangle\langle 0|)^{\otimes(n+n_A)} U_{\theta}^{\dagger}]. \tag{B2}$$

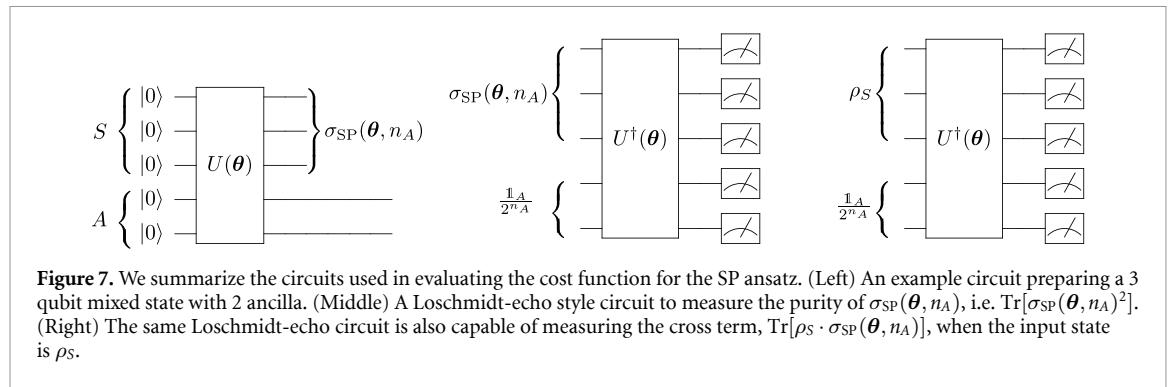


Figure 7. We summarize the circuits used in evaluating the cost function for the SP ansatz. (Left) An example circuit preparing a 3 qubit mixed state with 2 ancilla. (Middle) A Loschmidt-echo style circuit to measure the purity of $\sigma_{SP}(\theta, n_A)$, i.e. $\text{Tr}[\sigma_{SP}(\theta, n_A)^2]$. (Right) The same Loschmidt-echo circuit is also capable of measuring the cross term, $\text{Tr}[\rho_S \cdot \sigma_{SP}(\theta, n_A)]$, when the input state is ρ_S .

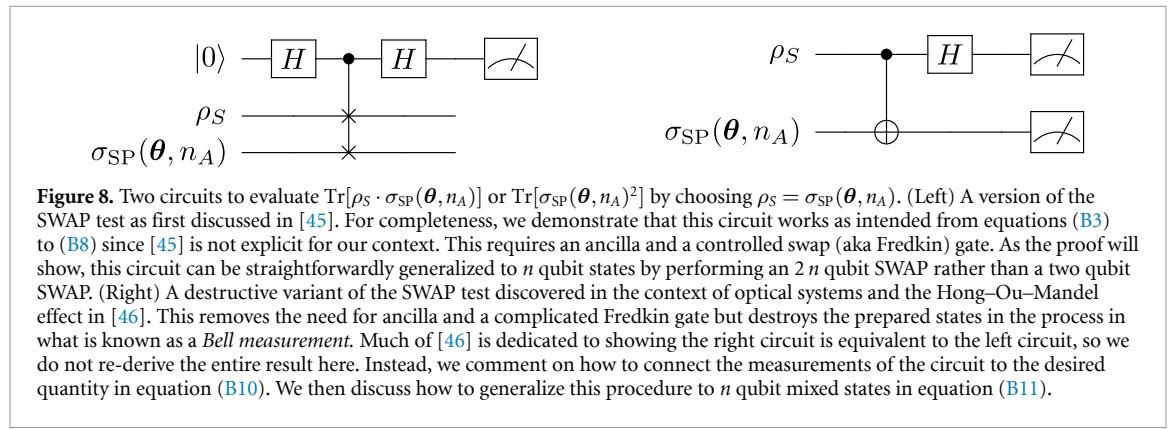


Figure 8. Two circuits to evaluate $\text{Tr}[\rho_S \cdot \sigma_{SP}(\theta, n_A)]$ or $\text{Tr}[\sigma_{SP}(\theta, n_A)^2]$ by choosing $\rho_S = \sigma_{SP}(\theta, n_A)$. (Left) A version of the SWAP test as first discussed in [45]. For completeness, we demonstrate that this circuit works as intended from equations (B3) to (B8) since [45] is not explicit for our context. This requires an ancilla and a controlled swap (aka Fredkin) gate. As the proof will show, this circuit can be straightforwardly generalized to n qubit states by performing an $2n$ qubit SWAP rather than a two qubit SWAP. (Right) A destructive variant of the SWAP test discovered in the context of optical systems and the Hong–Ou–Mandel effect in [46]. This removes the need for ancilla and a complicated Fredkin gate but destroys the prepared states in the process in what is known as a *Bell measurement*. Much of [46] is dedicated to showing the right circuit is equivalent to the left circuit, so we do not re-derive the entire result here. Instead, we comment on how to connect the measurements of the circuit to the desired quantity in equation (B10). We then discuss how to generalize this procedure to n qubit mixed states in equation (B11).

The translation of this procedure into a quantum circuit is straightforward and is shown for an $n = 3, n_A = 2$ example in the left-most circuit of figure 7. Given $U(\theta)$, we can measure an estimate of our cost function using the remaining two Loschmidt echo style circuits shown in figure 7. The middle circuit evaluate the purity term $\text{Tr}[\sigma_{SP}(\theta, n_A)^2]$ and the right-most circuit evaluates the cross term $\text{Tr}[\rho_S \cdot \sigma_{SP}(\theta, n_A)]$ which we showed in the main text equation (17). Note that we omitted providing a circuit to measure the purity of ρ_S since this is a static term during the optimization anyway. If desired, one could simply use middle circuit and replace $\sigma_{SP}(\theta, n_A)$ with ρ_S .

In the middle purity evaluation circuit, we avoid writing out the circuit which prepares $\sigma_{SP}(\theta, n_A)$ a second time. However, it is worth mentioning that generating $\sigma_{SP}(\theta, n_A)$ itself takes n_A ancilla, so the middle circuit takes a minimum of $n + 2n_A$ qubits without resetting the ancilla A to use twice. As for preparing the totally mixed state on the ancilla system, this can be done in two ways. In the first, we prepare one of 2^{n_A} basis states with a uniformly random probability for every shot that C is evaluated. Alternatively, one may choose to actually prepare the uniformly mixed state which can be done by preparing any completely entangled bi-partite state (i.e. a GHZ state) on $2n_A$ qubits. In the latter case, the middle circuit therefore uses $n + 3n_A$ qubits.

We also draw attention to the fact that we use $U(\theta)$ to prepare $\sigma_{SP}(\theta, n_A)$ in the first circuit but $U^\dagger(\theta)$ in the two cost function evaluation circuits. The use of $U^\dagger(\theta)$ is fleshed out mathematically in equation (17) of the main text. Intuitively, it is as if U^\dagger is undoing the preparation circuit U —hence the name Loschmidt-echo like circuit. In fact, the final cost term is ultimately evaluated by counting the number of 0s obtained at the final registers which corroborates this intuition.

As discussed in the main text, the Loschmidt echo circuits we cooked up in figure 7 are NISQ friendly but incur a poor shot scaling as n_A grows. If in practice, n_A is on the order of n and both are large, then it is best to instead use SWAP test circuits [45] or their destructive variant [46]. The SWAP and destructive SWAP circuits to evaluate $\text{Tr}[\rho_S \cdot \sigma_{SP}(\theta, n_A)]$ are shown in figure 8. Note that these also can be used for $\text{Tr}[\sigma_{SP}(\theta, n_A)^2]$ by replacing ρ_S with $\sigma_{SP}(\theta, n_A)$. While the references ultimately contain sufficient information to deduce that these circuits work as claimed, it is not obvious at a glance. For posterity and completeness, we provide a tailored derivation of the claim in the present context.

We begin by showing that the left SWAP test circuit is sufficient to measure $\text{Tr}[\rho\sigma]$ (we have dropped subscripts for simplicity). The first thing we need to know for this derivation is that

$$\text{Tr}[(\rho \otimes \sigma)\text{SWAP}] = \text{Tr}[\rho\sigma]. \quad (\text{B3})$$

As we will see, our goal will be to find an observable on the ancilla qubit alone whose expectation value gives us the left-hand side of equation (B3). From this identity, the desired outcome follows, and this is why we call it a SWAP test method.

Consider the state right after the controlled SWAP (aka Fredkin) gate which we shall call ω for concreteness. By simple Dirac notation manipulations, we arrive at,

$$\omega \equiv (|0\rangle\langle 0| \otimes \mathbb{1} + |1\rangle\langle 1| \otimes \text{SWAP}) (|+\rangle\langle +| \otimes \rho \otimes \sigma) (|0\rangle\langle 0| \otimes \mathbb{1} + |1\rangle\langle 1| \otimes \text{SWAP}) \quad (\text{B4})$$

$$= \frac{1}{2} \{ |0\rangle\langle 0| \otimes \rho \otimes \sigma + |0\rangle\langle 1| \otimes (\rho \otimes \sigma \text{ SWAP}) + |1\rangle\langle 0| \otimes (\text{SWAP} \rho \otimes \sigma) + |1\rangle\langle 1| \otimes \sigma \otimes \rho \}. \quad (\text{B5})$$

The second line can be thought of as block matrix,

$$\omega = \frac{1}{2} \begin{pmatrix} \rho \otimes \sigma & (\rho \otimes \sigma) \text{SWAP} \\ \text{SWAP}(\rho \otimes \sigma) & \sigma \otimes \rho \end{pmatrix}. \quad (\text{B6})$$

Right before measurement we are then left with the state $H\omega H$. Since we then only make a measurement on the ancilla qubit, we will only need the diagonal entries of this block matrix to choose the right observable to measure. In particular, we find

$$\begin{aligned} H\omega H \\ = \frac{1}{4} \begin{pmatrix} \rho \otimes \sigma + (\rho \otimes \sigma) \text{SWAP} + \text{SWAP}(\rho \otimes \sigma) + \sigma \otimes \rho & \dots \\ \dots & \rho \otimes \sigma - (\rho \otimes \sigma) \text{SWAP} - \text{SWAP}(\rho \otimes \sigma) + \sigma \otimes \rho \end{pmatrix}. \end{aligned} \quad (\text{B7})$$

Staring at equation (B7) in light of equation (B3), we see that the right observable is σ_z . Indeed,

$$\text{Tr}[\sigma_z H\omega H] = \text{Tr}[\rho \otimes \sigma \text{SWAP}] = \text{Tr}[\rho\sigma]. \quad (\text{B8})$$

Note that the above derivation really only relied on the fact that $\text{SWAP}(\rho \otimes \sigma) \text{SWAP} = \sigma \otimes \rho$. Hence, this circuit works to evaluate $\text{Tr}[\rho\sigma]$ when ρ and σ are composed of an arbitrary number of qubits despite the way our diagram implies they are single qubit states. We simply chose to write it this way since it is easier to interpret and reason in this way afterwards then try to worry about multiple qubits from the start.

Next, we discuss the right destructive SWAP variation. Let's again begin by considering that ρ and σ are single qubit states. It turns out, the relevant measurement is the projection into the $|11\rangle\langle 11|$ subspace. In particular, let

$$P_{11} = \text{Tr}[|11\rangle\langle 11| \cdot H \cdot \text{CNOT} \cdot \rho \otimes \sigma \cdot \text{CNOT} \cdot H] \quad (\text{B9})$$

be the probability that the final result from the destructive SWAP circuit is $|11\rangle$. Then,

$$\text{Tr}[\rho\sigma] = 2(1 - P_{11}) - 1 \quad (\text{B10})$$

gives us the desired quantity.

An n qubit generalization is not quite as straightforward as the SWAP test generalization. In particular, we replace the single CNOT and Hadamard with a transversal application of CNOTs and Hadamards—i.e. figure 11 in [46]. Further, we do not just project onto $|1\dots 1\rangle\langle 1\dots 1|$. In fact, the augmented procedure is actually easier to state at the level of individual measurements rather than projectors. Supposing ρ and σ are n qubit states, then we can label the measurement outcomes for the first n ρ registers with a bitstring \mathbf{a} and of the σ registers \mathbf{b} . We say the test ‘fails’ when the bit-wise and of the two bit-strings has odd parity, i.e. $|\mathbf{a} \wedge \mathbf{b}|$ is odd. Identifying P_f as this failure probability obtained by repeating this procedure ad infinitum, we again find

$$\text{Tr}[\rho\sigma] = 2(1 - P_f) - 1. \quad (\text{B11})$$

As a sanity check, we can confirm that P_f corresponds to P_{11} for the single qubit case. Here, the bit-strings each have one element, so the condition reduces to $|\mathbf{a} \wedge \mathbf{b}| = 1$ which occurs if and only if $a = b = 1$. An alternative derivation for the n-qubit destructive SWAP test viewed as a Bell basis measurement can be found in [90] (page 8) whose intuition can be understood from [91] (page 6).



Figure 9. We summarize the circuits used to evaluate the cost function for the CCPS ansatz. (Left) A generic way to prepare a randomly sampled eigenvector of the CCPS state. Here, the state $|i_1 i_2 \dots i_n\rangle$ is a short-hand for an n qubit computational basis state, so each i_j is either 0 or 1. Given the randomly sampled bitstring, we can prepare the desired state with single qubit X gates, and then we apply the same unitary U_θ regardless of the randomly sampled input. (Right) Here, we provide a circuit to sample the distribution $q_\theta(i)$ defined in equation (B14). We simply prepare ρ_S , apply U_θ^\dagger , and measure each qubit separately in the computational basis. By equation (B14) and the fact that $p_\phi(i)$ is stored classically, we can therefore compute the desired cross term using a *classical SWAP test* as described in the main text around equation (24).

B.1.2. CCPS ansatz circuits

A mixed state is often thought as a probabilistic mixture of pure states. The CCPS ansatz is a direct implementation of this idea,

$$\sigma_{\text{CCPS}}(\alpha, R) := \sum_{i=0}^{R-1} p_\phi(i) U_\theta |i\rangle\langle i| U_\theta^\dagger. \quad (\text{B12})$$

Namely, experimental observables are obtained by averaging over many experiments in which the input state is $U_\theta |i\rangle$ chosen with probability $p_\phi(i)$ ¹³. The ensemble of circuits can thus be represented by U_θ acting on an input basis vector $|i\rangle$ as in figure 9. One can choose any set of basis vector $\{|i\rangle\}_{i=0}^{R-1}$, but for our work, we choose the set of computational basis states¹⁴. Given a the preparation unitary U_θ , we can use the right circuit to sample the distribution

$$q_\theta(i) := \langle i | U_\theta^\dagger \rho U_\theta | i \rangle. \quad (\text{B13})$$

As discussed in the text surrounding equation (24), this distribution—alongside the classically stored $p_\phi(i)$ —allow us to compute the cross term

$$\text{Tr}[\rho\sigma] = \sum_i p_\phi(i) q_\theta(i) \quad (\text{B14})$$

by using the *classical SWAP test*.

As before, we omit a procedure to estimate $\text{Tr}[\rho_S^2]$ which does not affect the optimization. This time, however, we have also omitted a circuit to estimate $\text{Tr}[\sigma_{\text{CCPS}}^2]$ since this is done entirely classically again using the classical SWAP test. Other than that, the only subtlety in our circuits is simply that we have not explicitly written out different lines for different qubits. The reason is that for this ansatz, there are no ancilla qubits necessary, so it is understood that each line is for n qubits. As mentioned in the caption, the measurement is a single qubit measurement on each qubit.

B.2. The parameterized circuits used for $U(\theta)$

So far, we have described all the circuits where $U(\theta)$ was understood to be some PQC. Here, we define and justify the choices we make for the different classes of states we consider. We begin with an abstract description of a hardware efficient tiling. We then discuss how this tiling is applied for the parameterized circuits in the SP ansatz and the CCPS ansatz. Then we move into specifics for Bures random states, XY thermal states, and hardware noise induced states. As discussed in the main text, the CCPS ansatz automatically provides a description of the principal components of the target state ρ in the computational basis. This is not generally true for the SP ansatz, and we conclude with a discussion of how to generate an SP ansatz that does allow for extraction of the principal components but note that it is generally not practical.

¹³ When R is small, we can think of $p_\phi(i)$ as a probability vector with entries p_i stored in classical memory. The funny ϕ notation conveys the fact that R can be exponentially large in general, and more involved means of storing and sampling this distribution must be used in this case.

¹⁴ This is not just a matter of simplicity. The value in this choice is that we can reach any computational basis state with a depth 1 circuit—namely one applying an X gate on those qubits initialized in $|1\rangle$ and identity otherwise. Thus, to reach a ‘complex state,’ we must use a deep circuit ansatz U_θ . A different choice for $|i\rangle$ could lead to hiding the necessary complexity of the circuit ansatz in practice where $|0\rangle^{\otimes n}$ is the initial state.

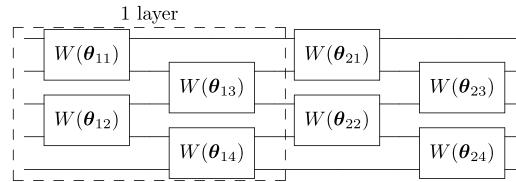


Figure 10. A hardware efficient tiling of a two-qubit gate $W(\theta)$ across a linearly connected five-qubit circuit. The notation θ_{lk} is meant to convey the parameters of the k th gate in the l th layer. Each layer consists of $n - 1$ applications of W for n qubits.

B.2.1. Hardware efficient tiling and its use in the CCPS and SP ansätze

Let $W(\theta)$ be an unspecified two-qubit gate parameterized by a vector of angles, θ . A hardware efficient tiling of W is given in figure 10. A single layer consists of what is shown in the ‘dotted rectangle’ on the left. The name hardware-efficient comes from the fact that for a linearly connected device, only neighboring qubits need to be coupled, and so no swaps are needed. Furthermore, each layer consists of a depth- $2D_W$ circuit only, where D_W is the depth of W itself.

For the CCPS ansatz, we simply apply a hardware efficient tiling on the n qubits. That is,

$$U_{\theta}^{\text{CCPS}} = \prod_{l=1}^L \prod_{k=1}^{n-1} W(\theta_{lk}), \quad (\text{B15})$$

where the product over k represents the application of each gate in layer l and the product over l is for L layers, which is chosen depending on the class of state. For the SP ansatz, we simply apply the same tiling but to the $n + n_A$ system plus ancilla qubits,

$$U_{\theta}^{\text{SP}} = \prod_{l=1}^L \prod_{k=1}^{n+n_A-1} W(\theta_{lk}). \quad (\text{B16})$$

B.2.2. Bures random state ansatz

There are no non-trivial symmetry operators S for which $[\rho, S] = 0$ for all ρ drawn from the Bures measure. Without any inherent symmetry structure, we choose the most generic W : an arbitrary two-qubit gate; i.e. W can express any rotation in $U(4)$. By the KAK decomposition [92–94], we can decompose W using three CNOTs and 15 elementary single qubit rotations (see figure 7 in [94] for example). Next, note that one way to generate ρ from the Bures distribution is through applying a Haar random unitary on $2n$ qubits, and then putting this state in coherent superposition with the same state changed by a local transformation on the n system qubits and then tracing out the ancilla [70] (see equation (D4)). This suggests that no low-depth circuit exists to faithfully generate ρ , so we choose $L = n$. This gives a total of $15n(n - 1)$ trainable parameters. As discussed in Cerezo *et al* [21], we expect a linear depth alternating ansatz with this many parameters to exhibit a barren plateau and thus to not be scalable. However, in fact, we know that trying to learn a Haar random unitary induces a barren plateau regardless of the choice of ansatz when no other information is known [95]. Hence, this class of states is likely not scalable beyond the small sizes testable in the NISQ era anyway. To that end, it serves as a proof of principle that even the most difficult states (for tractable sizes) can be learned by our method.

B.2.3. XY model ansatz

The XY model does exhibit symmetries. In particular, the spin model is particle conserving, and as a chain in 2D, it is invariant under any global rotation. As discussed in [96], the structure of the XY model can be used to design a generically good ansatz using $\mathcal{O}(n^2)$ gates with circuit depth $\mathcal{O}(n \log n)$. However, as discussed in appendix 3c of [97], one can instead use n alternating layers of Givens rotations which still obeys the symmetries but results in a lower $2n$ depth with a gate count of $n^2 - n$. The Givens rotation is a single parameter gate,

$$G(\theta) = \begin{pmatrix} 1 & 0 & 0 & 0 \\ 0 & \cos(\theta/2) & -\sin(\theta/2) & 0 \\ 0 & \sin(\theta/2) & \cos(\theta/2) & 0 \\ 0 & 0 & 0 & 1 \end{pmatrix}, \quad (\text{B17})$$

which rotates in the subspace where $|00\rangle$ and $|11\rangle$ are fixed.

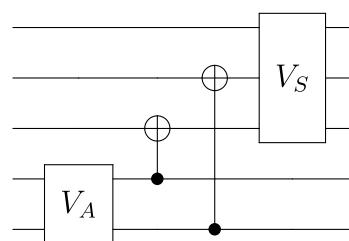


Figure 11. SP PCA Ansatz. An example SP PCA ansatz acting on $3 + 2$ qubits. We first apply a gate V_A on the $n_A = 2$ ancilla, connect the two sets of qubits with a CNOT cascade, and then apply a gate V_S on the system qubits.

In our work, however, we want to associate the state $|0\rangle^{\otimes n}$ in the CCPS ansatz to the first principal component of ρ . Hence, $|0\rangle^{\otimes n}$ should not in general be preserved. As a fix, we can consider other Givens rotations that differ from $G(\theta)$ by an arbitrary permutation of basis elements. For example,

$$G'(\theta) = \begin{pmatrix} 0 & 1 & 0 & 0 \\ \cos(\theta/2) & 0 & 0 & \sin(\theta/2) \\ -\sin(\theta/2) & 0 & 0 & \cos(\theta/2) \\ 0 & 0 & 1 & 0 \end{pmatrix}, \quad (\text{B18})$$

is also a valid Givens rotation where the set of fixed states is permuted. To avoid attaching to a fixed choice, we define a Givens gate as

$$W(\theta_1, \theta_2, \theta_3) = G(\theta_1)G'(\theta_2)G(\theta_3), \quad (\text{B19})$$

which is analogous to an Euler decomposition of a rotation in 3D as $R_X(\theta_1)R_Y(\theta_2)R_X(\theta_3)$. Using a tiling of this W , we find empirically that a depth $L = \log n$ is sufficient, and so use a total of $3(n-1)\log(n)$ single-parameter gates with depth $3\log(n)$, which represents yet another improvement over [42].

B.2.4. Ansatz for hardware implementation

To learn the purification (i.e. the SP ansatz) of the single qubit state ρ_+ , we used an arbitrary two-qubit gate built up with the previously mentioned KAK decomposition. For the pure state approximation, we simply forgo using an ancilla. In other words, for the $R = 1$ approximation of ρ_+ , we used a single qubit circuit ansatz,

$$W_{\rho_+}^{R=1}(\theta_1, \theta_2, \theta_3) = R_Z(\theta_1)R_Y(\theta_2)R_Z(\theta_3), \quad (\text{B20})$$

on the system qubit.

To learn the principal components and values of ρ_{Φ^+} with the CCPS ansatz, we use an arbitrary two-qubit circuit U_2 . The principal components are then obtained as $\{U_2|00\rangle, U_2|01\rangle, U_2|10\rangle, U_2|11\rangle\}$ whereas the principal values are stored as a vector we train over, $\{p_{00}, p_{01}, p_{10}, 1 - (p_{00} + p_{01} + p_{10})\}$.

B.3. SP PCA ansatz

We note that it is possible to construct an SP ansatz that allows for the principal components of ρ to be extracted using measurements in the computational basis. For clear reasons, we call this the PCA ansatz, and we show a generic example of such an ansatz in figure 11. The idea bears some similarities to the ansatz recently presented in [98].

To see why this is a PCA ansatz, we first consider the action of a generic version of the circuit in figure 11, proceeding step by step. The state before the CNOT cascade is as follows:

$$|\phi_1\rangle \equiv (\mathbb{1}_S \otimes V_A) |0\rangle_S^{\otimes n_S} |0\rangle_A^{\otimes n_A} \quad (\text{B21a})$$

$$= |0\rangle_S^{\otimes n_S} \sum_{i_1, i_2, \dots, i_{n_A}=0}^1 c_{i_1, i_2, \dots, i_{n_A}} |i_1, i_2, \dots, i_{n_A}\rangle_A. \quad (\text{B21b})$$

After applying the CNOTs, we get

$$|\phi_2\rangle \equiv \prod_{i=0}^{n_A-1} \text{CNOT}_{n_S+i, n_S-(i+1)} |\phi_1\rangle \quad (\text{B22a})$$

$$= |0\rangle_S^{\otimes n_S-n_A} \sum_{i_1, i_2, \dots, i_{n_A}=0}^1 c_{i_1, i_2, \dots, i_{n_A}} |i_1, i_2, \dots, i_{n_A}\rangle_S |i_1, i_2, \dots, i_{n_A}\rangle_A \quad (\text{B22b})$$

$$= \sum_{i_1, i_2, \dots, i_{n_A}=0}^1 c_{i_1, i_2, \dots, i_{n_A}} |0, \dots, 0, i_1, i_2, \dots, i_{n_A}\rangle_S |i_1, i_2, \dots, i_{n_A}\rangle_A \quad (\text{B22c})$$

$$= \sum_{k=0}^{2^{n_A}-1} c_k |\tilde{k}\rangle_S |k\rangle_A, \quad (\text{B22d})$$

where we introduced an arbitrary simpler indexing at the end. The choice of *tilde* on the *S* basis ket is to emphasize that there is a leading $n_S - n_A$ qubits in the all-zeros state.

By applying a system local unitary V_S and then tracing out the ancilla, we get

$$\sigma_S \equiv \text{Tr}_A \left[\sum_{k,j=0}^{2^{n_A}-1} c_k c_j^* (V_S \otimes \mathbb{1}_A) |\tilde{k}\rangle_S |k\rangle_A \langle j|_S \langle j|_A (V_S^\dagger \otimes \mathbb{1}_A) \right] \quad (\text{B23a})$$

$$= \sum_{k=0}^{2^{n_A}-1} p_k V_S |\tilde{k}\rangle \langle \tilde{k}| V_S^\dagger, \quad (\text{B23b})$$

where in the last step we simply identified $|c_k|^2 = p_k$ as probabilities. For arbitrary V_A and V_S , the state σ_S is an arbitrary rank- 2^{n_A} density matrix, and hence, this is clearly a legitimate ansatz for a density matrix on *S*.

At the same time, this ansatz prepares the state

$$\sigma_A \equiv \text{Tr}_S[|\phi_2\rangle \langle \phi_2|] = \sum_{k=1}^{2^{n_A}} p_k |k\rangle \langle k|_A \quad (\text{B24})$$

on the ancilla system. By inspection of σ_A and $|\phi_2\rangle$ we see that a measurement of the ancilla in the computational basis yields the state

$$|u_k\rangle \equiv V_S |\tilde{k}\rangle_S |k\rangle_A \quad (\text{B25})$$

with probability p_k . Hence, a measurement of the ancilla prepares the eigenstates of σ_S , $V_S |\tilde{k}\rangle$ with corresponding probability p_k . Provided $\sigma_S \approx \rho$ as a result of a successful training, this gives us a way to probabilistically prepare the first 2^{n_A} principal components of ρ . In particular, by repeatedly preparing σ_S and measuring the ancilla, we prepare the k th principal component with probability p_k .

Appendix C. Gradient analysis for SP ansatz

The HS distance cost for the SP ansatz takes the form $C_{\text{SP}(\theta, n_A)} = \text{Tr}[\rho^2] + \text{Tr}[\sigma(\theta, n_A)^2] - 2\text{Tr}[\rho\sigma(\theta, n_A)]$, where the trial state $\sigma(\theta, n_A)$ is found via its purification, i.e. $\sigma_{\text{SP}}(\theta, n_A) := \text{Tr}_A[U_\theta(|0\rangle\langle 0|)^{\otimes(n+n_A)} U_\theta^\dagger]$, for our vector of training parameters θ . The gradient with respect to θ is given by

$$\nabla C_{\text{SP}(\theta, n_A)} = \nabla \text{Tr}[\rho^2] + \nabla \text{Tr}[\sigma(\theta, n_A)^2] - 2\nabla \text{Tr}[\rho\sigma(\theta, n_A)]. \quad (\text{C1})$$

The $\text{Tr}[\rho^2]$ term vanishes trivially due to a lack of dependence on θ , whilst the overlap term $\text{Tr}[\rho\sigma(\theta, n_A)]$ obeys the standard parameter shift rule [57, 58]. In our work, we specifically use the ‘Pauli parameter shift rule’ in which the circuit ansätze are described using single qubit gates of the form $P_k = e^{-i\theta\sigma_k/2}$ for Pauli σ_k and parameter free CNOTs. In this case, the shift refers to running the circuit with $\theta \rightarrow \theta \pm \pi/2$ (this will become clear in the next equation). This choice is motivated by our use of IBM devices where this is the appropriate gate set.

We can still use the Pauli parameter shift rule for the $\text{Tr}[\sigma(\boldsymbol{\theta}, n_A)^2]$ term after applying a matrix differentiation rule [99, equation (11.175)]. Namely, given some function $f(x)$ we have that $\frac{\partial}{\partial \theta} \text{Tr}[f(A(\theta))] = \text{Tr} \left[g(A(\theta)) \frac{\partial A}{\partial \theta} \right]$ where $g(x) := \frac{\partial f}{\partial x}$. Thus we have that

$$\begin{aligned} \frac{\partial \text{Tr} [\sigma(\boldsymbol{\theta}, n_A)^2]}{\partial \theta_k} &= 2 \text{Tr} \left[\sigma(\boldsymbol{\theta}, n_A) \frac{\partial \sigma(\boldsymbol{\theta}, n_A)}{\partial \theta_k} \right] \\ &= \text{Tr} \left[\sigma(\boldsymbol{\theta}, n_A) \sigma \left(\boldsymbol{\theta} + \frac{\pi}{2} \cdot \mathbf{e}_k, n_A \right) \right] - \text{Tr} \left[\sigma(\boldsymbol{\theta}, n_A) \sigma \left(\boldsymbol{\theta} - \frac{\pi}{2} \cdot \mathbf{e}_k, n_A \right) \right], \end{aligned} \quad (\text{C2})$$

where in the second line we use the Pauli parameter shift rule (hence the $\pm \pi/2$) as applied directly to an operator rather than an expectation value. Thus we have the complete analytic expression to compute $\nabla C_{\text{SP}(\boldsymbol{\theta}, n_A)}$. We note that the gradient analysis for the local costs is entirely analogous.

Appendix D. Description of state ensembles

D.1. Bures random states

Before diving into what a Bures random state is specifically, it is helpful to discuss a few general considerations regarding generating random quantum states, as is done in the brief introduction of [69]. Generally, random quantum states are selected from a distribution that is invariant under global unitary transformations. For pure states, this single property uniquely defines a probability measure known as the *Haar measure*. But for mixed states, this property does not uniquely specify the measure. Assuming the distribution of eigenvectors and eigenvalues is independent, then we can write the probability measure over mixed states as [70]

$$d\mu = d\nu(\lambda_1, \lambda_2, \dots, \lambda_N) \times d\mu_V, \quad (\text{D1})$$

where $d\mu_V$ is the Haar measure and $d\nu$ is the measure over the normalized eigenvalues. So the problem is that we must also specify $d\nu$, which, while trivial for pure states, is unclear for mixed states.

There are many protocols one can follow to determine $d\nu$, but the most mathematically straightforward is to define a measure from the normalized volume elements of a metric [69, 70]. In this formalism, the key is then to choose an appropriate metric. The Bures distance, given by

$$D_B(\rho, \sigma) := \sqrt{2 - 2 \text{Tr} \left[\sqrt{\sqrt{\rho} \sigma \sqrt{\rho}} \right]}, \quad (\text{D2})$$

has many attractive properties as an unbiased choice. Many of these properties were first pointed out by Bures himself [100], but simpler explanations are provided in Hall [69] and Zyczkowski *et al* [70]. The summary in Zyczkowski *et al* [70] is especially concise: the Bures metric (i) has an interpretation as a distinguishability measure [101, 102], (ii) is the minimal monotone metric under quantum channels [103], and (iii) gives the statistical distance when applied to two diagonal operators [70]. The form of $d\nu$ is given in equation (13) of Zyczkowski *et al* [70], and for brevity, we call states drawn from this measure *Bures random states*.

For numerically tractable system sizes, generating a Bures random state is straightforward [70]. First, we generate a $2^n \times 2^n$ Ginibre random matrix G [104] with complex entries¹⁵. Then, we generate a Haar random unitary matrix U with the same dimensions. With these two matrices, the random state is given by

$$\rho_B = \frac{(\mathbb{1} + U)GG^\dagger(\mathbb{1} + U^\dagger)}{\text{Tr}[(\mathbb{1} + U)GG^\dagger(\mathbb{1} + U^\dagger)]}. \quad (\text{D3})$$

An alternative, more physically motivated way to generate the states is to construct a superposition of a random bipartite state $|\psi_1\rangle = U_{AB}|0,0\rangle$ with a local transformation of the same state, $|\psi_2\rangle = (V_A \otimes \mathbb{1})|\psi_1\rangle$ and then trace out the B degrees of freedom (see figure 5 in [70]),

$$\rho_A = \frac{\text{Tr}_B |\phi\rangle\langle\phi|}{\langle\phi|\phi\rangle} \quad |\phi\rangle \equiv [(\mathbb{1} + V_A) \otimes \mathbb{1}]|\psi_1\rangle. \quad (\text{D4})$$

¹⁵ The exact description is just a matrix whose entries are $G_{ij} = x + iy$ where $x, y \sim \mathcal{N}(0, 1)$.

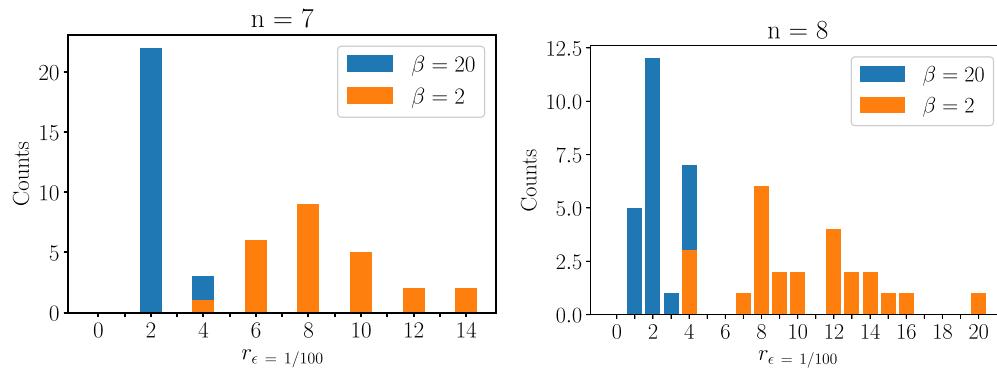


Figure 12. XY thermal state ϵ -rank depends on β . We show the distribution of ϵ -ranks (r_ϵ) for the $n = 7$ and $n = 8$ XY thermal state instances (25 per (n, β)). At $\beta = 20$ (low temperature), the resulting rank never goes beyond $r_\epsilon = 4$. For $n = 7$, the ground state is doubly degenerate, so $r_\epsilon \geq 2$ as well. For $n = 8$, the ground state is unique, so $r_\epsilon = 1$ is reached for some of the states. For $\beta = 2$, the ranks go from $r_\epsilon = 4$ all the way up to $r_\epsilon = 14, 20$ for $n = 7, 8$, respectively.

D.2. XY thermal states

By XY thermal states, we mean states of the form

$$\rho_n^{(XY)} = \frac{e^{-\beta H_{XY}}}{\text{Tr}[e^{-\beta H_{XY}}]}, \quad (\text{D5a})$$

$$H_{XY} = \sum_{i=1}^{n-1} J_i X_i X_{i+1} + K_i Y_i Y_{i+1}, \quad (\text{D5b})$$

where $J_i, K_i \sim \mathcal{N}(0, 1)$ are i.i.d. normal random variables. By controlling β , we control the effective rank r_ϵ of the random states generated in this way, which we demonstrate in figure 12 with $\epsilon = 1/100$. The reason is simple: At $\beta = \infty$ (i.e. $T = 0$), we expect a system with a non-degenerate ground state to be in its pure ground state which has $r_\epsilon = 1$. For a degenerate ground state, we get a totally mixed state in the ground-state subspace, so if the degeneracy is g_0 , we find $r_\epsilon = g_0$. For small but non-zero temperatures ($\beta = 20$ here), the state is a convex combination of low-lying energy states, so $r_\epsilon > g_0$ for most choices of ϵ . As β decreases, the state becomes closer to the totally mixed state until it actually reaches it at $\beta = 0$. By choosing a larger intermediate temperature ($\beta = 2$ here), we interpolate between the two extremes and simply get a Gibbs state with large ϵ -rank but without being completely mixed.

Despite the random states having different coefficients (and subsequent rank), the model still respects important symmetries. Most importantly, our XY model is a chain in 2D. Hence, it is invariant under any global rotation of all spins, and furthermore, the number of spins is constant. These symmetries allow us to greatly simplify the choice of ansatz to one that respects these symmetries, as discussed in appendix B.

Appendix E. Computing figure of merit

Our figure of merit comes from the solution to the QLRAP [18]. Namely, the QLRAP is to find the state $\sigma_{\text{opt}}(R)$ that satisfies

$$\sigma_{\text{opt}}(R) = \arg \min_{\sigma \geq 0, \text{rank}(\sigma) \leq R, \text{Tr}(\sigma) = 1} D(\rho, \sigma)^2. \quad (\text{E1})$$

As shown in [18], the unique optimal solution for the HS distance is given by

$$\sigma_{\text{opt}}(R) = \tau_R + N_R, \quad (\text{E2a})$$

$$\tau_R \equiv \Pi_R \rho \Pi_R, \quad (\text{E2b})$$

$$N_R \equiv \frac{1 - \text{Tr}[\tau_R]}{R} \Pi_R, \quad (\text{E2c})$$

with corresponding minimal HS cost

$$C(\alpha_{\text{opt}}, R) = \text{Tr}[(I - \Pi_R) \rho^2] + \text{Tr}[N_R^2]. \quad (\text{E3})$$

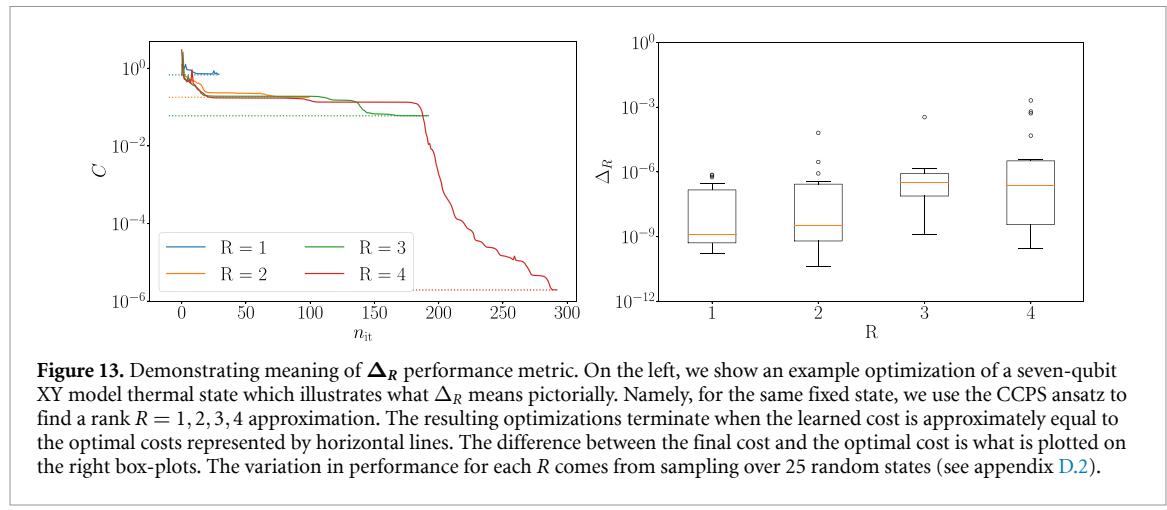


Figure 13. Demonstrating meaning of Δ_R performance metric. On the left, we show an example optimization of a seven-qubit XY model thermal state which illustrates what Δ_R means pictorially. Namely, for the same fixed state, we use the CCPS ansatz to find a rank $R = 1, 2, 3, 4$ approximation. The resulting optimizations terminate when the learned cost is approximately equal to the optimal costs represented by horizontal lines. The difference between the final cost and the optimal cost is what is plotted on the right box-plots. The variation in performance for each R comes from sampling over 25 random states (see appendix D.2).

With a little algebra, it is easy to see that the first term corresponds to the sum of squares of the $r - R$ eigenvalues of ρ not approximated (assuming $\text{rank}(\rho) = r$), and the second term accounts for the constant offset between the first R eigenvalues of ρ and the R re-normalized eigenvalues of σ_{opt} .

Any compilation of ρ as described throughout our paper will find an empirical cost $C(\alpha^*, R) \geq C(\alpha_{\text{opt}}, R)$, so a simple figure of merit is just their difference:

$$\Delta_R \equiv C(\alpha^*, R) - C(\alpha_{\text{opt}}, R). \quad (\text{E4})$$

In figure 13, we help clarify how we compute this difference for an example optimization of a seven-qubit XY thermal state using the CCPS ansatz with $R \in \{1, 2, 3, 4\}$. Pictorially, Δ_R is the vertical distance between the final cost (solid line) and the optimal cost (horizontal dashed line).

Appendix F. Additional hardware results

In section 4, we successfully compiled full rank approximations of the one qubit states ρ_{HS} and $\tilde{\rho}_+$ and the two qubit state $\tilde{\rho}_{\Phi+}$ (see the referenced section for state definitions). In particular, within an allotted budget of 100 Powell iterations, we optimized to a cost value on the order of $10^{-3} \leq C_{\text{noiseless}} \leq 10^{-2}$ which is close to the precision floor allowed by our use of 10^5 shots to evaluate the cost function. Here, we present additional hardware results where we compile lower rank approximations of these states either directly as in figure 14 or indirectly through truncation of the learned full rank CCPS state as in figure 15. For convenience, we have included the full-rank optimizations shown in section 4 alongside the lower rank optimizations.

In figure 14(a), we compile a full rank (i.e. rank two using $n_A = 1$ ancilla) and a pure state ($n_A = 0$) approximation of the single qubit state $\tilde{\rho}_+$ using the SP ansatz. In the full rank case, we find $C_{\text{noiseless}}(n_A = 1) = \Delta_2 \approx 6.44 \times 10^{-3}$. In the pure state case, we find $C_{\text{noiseless}}(n_A = 0) = 2.5 \times 10^{-2}$, but the optimal possible cost is $C^*(n_A = 0) = 4.1 \times 10^{-3}$. Thus, $\Delta_1 = 2.1 \times 10^{-2}$, which is an acceptable final value on the order of 10^{-2} . It is also interesting to note that around $n_{\text{it}} = 10$, the difference is minimal, reaching $C_{\text{noiseless}}(n_A = 0) = 4.3 \times 10^{-3} \implies \Delta_1 = 2.0 \times 10^{-3}$. However, it is not forthright to report this as the final found value since it relies on knowledge of $C_{\text{noiseless}}$ to pick the right n_{it} whereas our optimization stopping condition does not (i.e. it relies on C_{shot} only).

In figure 14(b), we compile a full rank ($R = 2$) and a pure state ($R = 1$) approximation of the single qubit state ρ_{HS} using the CCPS ansatz. In the full rank case, we find $C_{\text{noiseless}}(R = 2) = \Delta_2 = 8.2 \times 10^{-3}$. In the pure state case, we find $C_{\text{noiseless}}(R = 1) = 4.3 \times 10^{-2}$, but the optimal possible cost is $C^*(R = 1) = 3.3 \times 10^{-2}$, so $\Delta_1 = 9.9 \times 10^{-3}$ which is an acceptable final value on the order of at least 10^{-2} .

Overall, we find that direct pure state compilations of ρ for both ansätze are learned to an acceptable value of Δ_1 . In addition, we find that finding this value takes fewer iterations for both ansätze suggesting that finding a lower-rank approximation is easier. This observations was also found for the Bures random and XY random thermal states in idealized noiseless, infinite shot classical simulations. Given that the lower-rank optimization requires fewer learnable parameters alongside these empirical results, we suspect this to be a general property of our algorithm. Namely, learning a lower rank approximation is easier.

In figure 15 we explore the quality of indirect compilations of $\tilde{\rho}_{\Phi+}$ by truncating a full-rank compilation. The results here mirror those in section 3.4 but for the hardware optimization of $\tilde{\rho}_{\Phi+}$. The story here is a bit more interesting on account of finite shot noise, however. The punchline is that due to having only a precision of roughly 10^{-2} , our $R = 4$ optimization found a local minimum, $C_{\text{noiseless}}(R = 4) = 6.8 \times 10^{-2}$

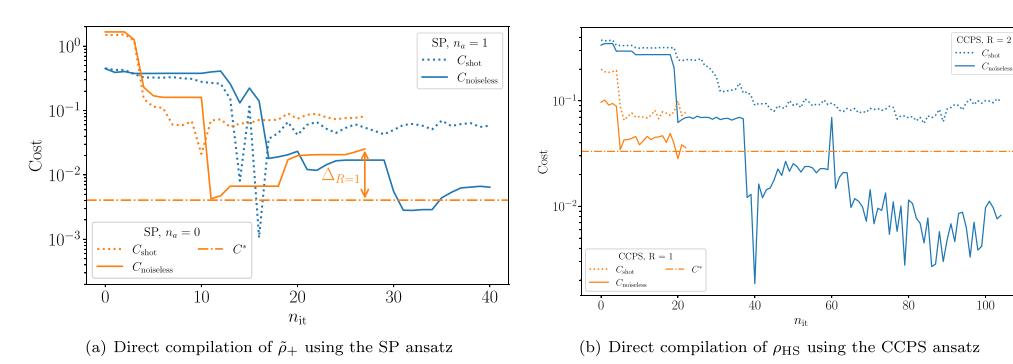


Figure 14. Direct low rank compilations of single qubit hardware states. We perform a direct full-rank and low-rank compilation of ρ_{HS} with the SP ansatz on the (left) and of $\tilde{\rho}_+$ with the CCPS ansatz on the (right). By direct we mean we performed a separate and independent compilation of the same state with an ansatz supporting a rank two or rank one approximation. Since these are one qubit states, this encompasses all interesting possibilities. For the rank one approximation, we also include a horizontal line, C^* , which denotes the lowest possible value of the noiseless cost. The difference between the final value of $C_{\text{noiseless}}$ and this line denotes Δ_R as shown explicitly on the (left) plot.

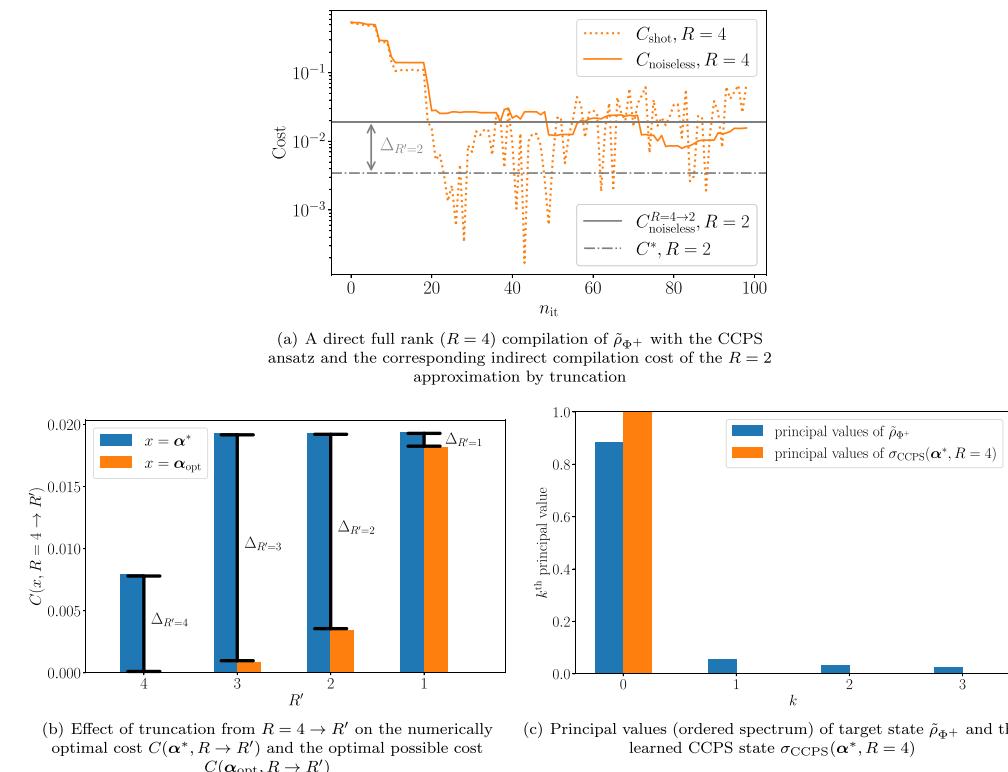


Figure 15. PCA analysis of $\rho_{\Phi+}$ when compiled on hardware using CCPS ansatz. In (a), we show the optimization result when finding a direct full-rank ($R = 4$) compilation of $\tilde{\rho}_{\Phi+}$. Given this CCPS ansatz, we can obtain any $R' < 4$ approximation for free by truncation. As an example, we plot the found cost when truncating to a rank two approximation, $C_{\text{noiseless}}^{R=4 \rightarrow 2}$, as a solid horizontal line. The optimal possible value for this cost, C^* , is shown in a dot-dashed line. The vertical distance between these lines denotes their difference, $\Delta R' = 2$. In (b), we show the quality of our truncation state for different truncation levels from $R' = 4$ (no truncation) to $R' = 1$ (maximum truncation to a pure state). For each R' , we plot the value of the found truncated cost $C(\alpha^*, R \rightarrow R')$ (left, blue) as well as the optimal possible cost $C(\alpha_{\text{opt}}, R \rightarrow R')$ (right, orange). The difference between these costs is $\Delta R'$ and is pictorially the height difference in the bars. In (c), we show the principal values (ordered eigenvalues) of the target state $\tilde{\rho}_{\Phi+}$ and the learned CCPS state $\sigma_{CCPS}(\alpha^*, R = 4)$.

(see figure 15(a)), that is effectively a rank one approximation. In other words, a pure state approximation of $\tilde{\rho}_{\Phi+}$ is sufficient to reach the cost noise floor, and this pure state solution was found in our hardware optimization.

With the punchline stated, the empirical quality of the truncated states is summarized in figure 15(b). Evidently, the pure state approximation ($R' = 1$) is closest to its optimal value since $\Delta R' = 1 = 1.2 \times 10^{-3}$. On the other hand, $\Delta R' > 1$ is an order of magnitude higher in all cases, i.e. $\Delta R' = 2 = 1.6 \times 10^{-2}$, $\Delta R' = 3 = 1.8 \times 10^{-2}$, and $\Delta R' = 4 = 6.8 \times 10^{-2}$. In fact, the quality monotonically decreases with R'

increasing. This suggests that the performance of our $R = 4$ compilation is dominated by the quality of the estimate of the first principal component. In figure 15(c), we verify this intuition by seeing that the numerically optimal CCPS state, $\sigma_{\text{CCPS}}(\alpha^*, R = 4)$, is essentially a pure state. Namely, it only has non-trivial support on the largest principal value. This interesting observation aside, we ultimately find that truncated costs have a comparable $\Delta_{R'}$ to the original Δ_R , so the indirect compilation of the full-rank CCPS state into lower rank approximations works well.

ORCID iDs

Nic Ezzell  <https://orcid.org/0000-0002-4983-0617>
 Elliott M Ball  <https://orcid.org/0000-0002-9060-9797>
 Aliza U Siddiqui  <https://orcid.org/0000-0001-8011-7863>
 Mark M Wilde  <https://orcid.org/0000-0002-3916-4462>
 Andrew T Sornborger  <https://orcid.org/0000-0001-8036-6624>
 Patrick J Coles  <https://orcid.org/0000-0001-9879-8425>

References

- [1] Nielsen M A and Chuang I L 2000 *Quantum Computation and Quantum Information* (Cambridge: Cambridge University Press)
- [2] Haah J, Harrow A W, Ji Z, Wu X and Yu N 2017 Sample-optimal tomography of quantum states *IEEE Trans. Inf. Theory* **63** 5628–41
- [3] O'Donnell R and Wright J 2016 Efficient quantum tomography *Proc. 48th Annual ACM Symp. on Theory of Computing* pp 899–912
- [4] Yuen H 2022 An improved sample complexity lower bound for quantum state tomography (arXiv:2206.11185)
- [5] Aaronson S 2019 Shadow tomography of quantum states *SIAM J. Comput.* **49** STOC18–368
- [6] Huang H-Y, Kueng R and Preskill J 2020 Predicting many properties of a quantum system from very few measurements *Nat. Phys.* **16** 1050–7
- [7] Elben A, Flammia S T, Huang H-Y, Kueng R, Preskill J, Vermersch B and Zoller P 2022 The randomized measurement toolbox *Nat. Rev. Phys.* **5** 9–24
- [8] Huang H-Y, Kueng R, Torlai G, Albert V V and Preskill J 2022 Provably efficient machine learning for quantum many-body problems *Science* **377** eabk3333
- [9] Paini M, Kalev A, Padilha D and Ruck B 2021 Estimating expectation values using approximate quantum states *Quantum* **5** 413
- [10] Khatri S, LaRose R, Poremba A, Cincio L, Sornborger A T and Coles P J 2019 Quantum-assisted quantum compiling *Quantum* **3** 140
- [11] Jones T and Benjamin S C 2022 Robust quantum compilation and circuit optimisation via energy minimisation *Quantum* **6** 628
- [12] Sharma K, Khatri S, Cerezo M and Coles P J 2020 Noise resilience of variational quantum compiling *New J. Phys.* **22** 043006
- [13] Verdon G, Marks J, Nanda S, Leichenauer S and Hidary J 2019 QuantumHamiltonian-based models and the variational quantum thermalizer algorithm (arXiv:1910.02071)
- [14] LaRose R, Tikku A, O'Neil-Judy É, Cincio L and Coles P J 2019 Variational quantum state diagonalization *npj Quantum Inf.* **5** 1–10
- [15] Cerezo M, Sharma K, Arrasmith A and Coles P J 2022 Variational quantum state eigensolver *npj Quantum Inf.* **8** 1–11
- [16] Schumacher B 1995 Quantum coding *Phys. Rev. A* **51** 2738
- [17] Cleve R and DiVincenzo D P 1996 Schumacher's quantum data compression as a quantum computation *Phys. Rev. A* **54** 2636
- [18] Ezzell N, Holmes Z and Coles P J 2022 The quantum low-rank approximation problem (arXiv:2203.00811)
- [19] Lloyd S and Weedbrook C 2018 Quantum generative adversarial learning *Phys. Rev. Lett.* **121** 040502
- [20] McClean J R, Boixo S, Smelyanskiy V N, Babbush R and Neven H 2018 Barren plateaus in quantum neural network training landscapes *Nat. Commun.* **9** 1–6
- [21] Cerezo M, Sone A, Volkoff T, Cincio L and Coles P J 2021 Cost function dependent barren plateaus in shallow parameterized quantum circuits *Nat. Commun.* **12** 1–12
- [22] Holmes Z, Arrasmith A, Yan B, Coles P J, Albrecht A and Sornborger A T 2021a Barren plateaus preclude learning scramblers *Phys. Rev. Lett.* **126** 190501
- [23] Holmes Z, Sharma K, Cerezo M and Coles P J 2022 Connecting ansatz expressibility to gradient magnitudes and barren plateaus *PRX Quantum* **3** 010313
- [24] Arrasmith A, Holmes Z, Cerezo M and Coles P J 2022 Equivalence of quantum barren plateaus to cost concentration and narrow gorges *Quantum Sci. Technol.* **7** 045015
- [25] Larocca M, Czarnik P, Sharma K, Muraleedharan G, Coles P J and Cerezo M 2022 Diagnosing barren plateaus with tools from quantum optimal control *Quantum* **6** 824
- [26] Sharma K, Cerezo M, Cincio L and Coles P J 2022 Trainability of dissipative perceptron-based quantum neural networks *Phys. Rev. Lett.* **128** 180505
- [27] Patti T L, Najafi K, Gao X and Yelin S F 2021 Entanglement devised barren plateau mitigation *Phys. Rev. Res.* **3** 033090
- [28] Thanaisilp S, Wang S, Nghiem N A, Coles P J and Cerezo M 2021 Subtleties in the trainability of quantum machine learning models (arXiv:2110.14753)
- [29] Uvarov A V and Biamonte J D 2021 On barren plateaus and cost function locality in variational quantum algorithms *J. Phys. A: Math. Theor.* **54** 245301
- [30] Ortiz Marrero C, Kieferová M and Wiebe N 2021 Entanglement-induced barren plateaus *PRX Quantum* **2** 040316
- [31] Cerezo M, Poremba A, Cincio L and Coles P J 2020 Variational quantum fidelity estimation *Quantum* **4** 248
- [32] Renner R and Wolf S 2004 Smooth rényi entropy and applications *Int. Symp. On Information Theory, 2004. ISIT 2004. Proc. (IEEE)* p 233
- [33] Renner R 2005 Security of quantum key distribution *PhD Thesis* School ETH Zurich (arXiv:quant-ph/0512258)

[34] Ohya M and Petz D 2004 *Quantum Entropy and Its Use* (Heidelberg: Springer) 31 Theorem 1.18

[35] Lloyd S, Mohseni M and Rebentrost P 2014 Quantum principal component analysis *Nat. Phys.* **10** 631–3

[36] Bravo-Prieto C, LaRose R, Cerezo M, Subasi Y, Cincio L and Coles P 2019 Variational quantum linear solver (arXiv:1909.05820)

[37] Knill E and Laflamme R 1998 Power of one bit of quantum information *Phys. Rev. Lett.* **81** 5672

[38] Fujii K, Kobayashi H, Morimae T, Nishimura H, Tamate S and Tani S 2018 Impossibility of classically simulating one-clean-qubit model with multiplicative error *Phys. Rev. Lett.* **120** 200502

[39] Morimae T 2017 Hardness of classically sampling the one-clean-qubit model with constant total variation distance error *Phys. Rev. A* **96** 040302

[40] Coles P J, Cerezo M and Cincio L 2019 Strong bound between trace distance and Hilbert-Schmidt distance for low-rank states *Phys. Rev. A* **100** 022103

[41] Gard B T, Zhu L, Barron G S, Mayhall N J, Economou S E and Barnes E 2020 Efficient symmetry-preserving state preparation circuits for the variational quantum eigensolver algorithm *npj Quantum Inf.* **6** 1–9

[42] Gibbs J, Gili K, Holmes Ze, Commeau B, Arrasmith A, Cincio L, Coles P J and Sornborger A 2021 Long-time simulations with high fidelity on quantum hardware (arXiv:2102.04313)

[43] Larocca M, Sauvage F, Sbahi F M, Verdon G, Coles P J and Cerezo M 2022 Group-invariant quantum machine learning *PRX Quantum* **3** 030341

[44] Bilkis M, Cerezo M, Verdon G, Coles P J and Cincio L 2021 A semi-agnostic ansatz with variable structure for quantum machine learning (arXiv:2103.06712)

[45] Barenco A, Berthiaume A, Deutsch D, Ekert A, Jozsa R and Macchiavello C 1997 Stabilization of quantum computations by symmetrization *SIAM J. Comput.* **26** 1541–57

[46] Carlos Garcia-Escartin J C and Chamorro-Posada P 2013 Swap test and hong-ou-mandel effect are equivalent *Phys. Rev. A* **87** 052330

[47] Yang Y 2022 Learning distributions by generative adversarial networks: approximation and generalization (arXiv:2205.12601)

[48] Montúfar G 2016 Restricted Boltzmann machines: introduction and review *Information Geometry and Its Applications IV* (Berlin: Springer) pp 75–115

[49] Subaş Y, Cincio L and Coles P J 2019 Entanglement spectroscopy with a depth-two quantum circuit *J. Phys. A: Math. Theor.* **52** 044001

[50] Guo X-Y et al 2023 Variational quantum simulation of thermal statistical states on a superconducting quantum processor *Chin. Phys. B* **32** 010307

[51] Vedral V and Plenio M B 1998 Entanglement measures and purification procedures *Phys. Rev. A* **57** 1619–33

[52] Schumacher B 1996 Sending entanglement through noisy quantum channels *Phys. Rev. A* **54** 2614–28

[53] Wilde M M 2013 *Quantum Information Theory* (Cambridge: Cambridge University Press) (<https://doi.org/10.1017/CBO9781139525343>)

[54] Chen R, Song Z, Zhao X and Wang X 2021 Variational quantum algorithms for trace distance and fidelity estimation *Quantum Sci. Technol.* **7** 015019

[55] Agarwal R, Rethinasamy S, Sharma K and Wilde M M 2021 Estimating distinguishability measures on quantum computers (arXiv:2108.08406)

[56] LaBorde M L, Rethinasamy S and Wilde M M 2021 Testing symmetry on quantum computers (arXiv:2105.12758)

[57] Mitarai K, Negoro M, Kitagawa M and Fujii K 2018 Quantum circuit learning *Phys. Rev. A* **98** 032309

[58] Schuld M, Bergholm V, Gogolin C, Izaac J and Killoran N 2019 Evaluating analytic gradients on quantum hardware *Phys. Rev. A* **99** 032331

[59] Ezzell N 2022 naezzell/qmsc: qmsc arXiv v1.0.1 release (Zenodo) (<https://doi.org/10.5281/zenodo.7039676>)

[60] Powell M J D 1964 An efficient method for finding the minimum of a function of several variables without calculating derivatives *Comput. J.* **7** 155–62

[61] Virtanen P et al 2020 Scipy 1.0: fundamental algorithms for scientific computing in python *Nat. Methods* **17** 261–72

[62] Nocedal J and Wright S 2006 *Numerical Optimization* 2nd edn (New York: Springer) 135–43

[63] Nelder J A and Mead R 1965 A simplex method for function minimization *Comput. J.* **7** 308–13

[64] Kraft D 1988 *A Software Package for Sequential Quadratic Programming* DFVLR-FB 88-28 German Aerospace Center — Institute for Flight Mechanics

[65] Powell M J D 1994 A direct search optimization method that models the objective and constraint functions by linear interpolation *Advances in Optimization and Numerical Analysis* (Berlin: Springer) pp 51–67

[66] Spall J C 1998 An overview of the simultaneous perturbation method for efficient optimization *Johns Hopkins APL Tech. Dig.* **19** 482–92

[67] Kübler J M, Arrasmith A, Cincio L and Coles P J 2020 An adaptive optimizer for measurement-frugal variational algorithms *Quantum* **4** 263

[68] Arrasmith A, Cincio L, Somma R D and Coles P J 2020 Operator sampling for shot-frugal optimization in variational algorithms (arXiv:2004.06252)

[69] Hall M J W 1998 Random quantum correlations and density operator distributions *Phys. Lett. A* **242** 123–9

[70] Życzkowski K, Penson K A, Nechita I and Collins B 2011 Generating random density matrices *J. Math. Phys.* **52** 062201

[71] Anschuetz E R and Kiani B T 2022 Beyond barren plateaus: quantum variational algorithms are swamped with traps (arXiv:2205.05786)

[72] Bittel L and Kliesch M 2021 Training variational quantum algorithms is NP-Hard *Phys. Rev. Lett.* **127** 120502

[73] Gambetta J M et al 2012 Characterization of addressability by simultaneous randomized benchmarking *Phys. Rev. Lett.* **109** 240504

[74] Mundada P, Zhang G, Hazard T and Houck A 2019 Suppression of qubit crosstalk in a tunable coupling superconducting circuit *Phys. Rev. Appl.* **12** 054023

[75] Tripathi V, Chen H, Khezri M, Yip K-W, Levenson-Falk E M and Lidar D A 2021 Suppression of crosstalk in superconducting qubits using dynamical decoupling (arXiv:2108.04530)

[76] Sheldon S, Magesan E, Chow J M and Gambetta J M 2016 Procedure for systematically tuning up cross-talk in the cross-resonance gate *Phys. Rev. A* **93** 060302

[77] Jozsa R, Horodecki M, Horodecki P and Horodecki R 1998 Universal quantum information compression *Phys. Rev. Lett.* **81** 1714

[78] Schumacher B and Westmoreland M D 2001 Indeterminate-length quantum coding *Phys. Rev. A* **64** 042304

[79] Horodecki M, Oppenheim J and Winter A 2007 Quantum state merging and negative information *Commun. Math. Phys.* **269** 107–36

[80] Plesch M and Bužek V 2010 Efficient compression of quantum information *Phys. Rev. A* **81** 032317

[81] Abdelhadi D and Renes J M 2020 On the second-order asymptotics of the partially smoothed conditional min-entropy & application to quantum compression *IEEE J. Sel. Areas Inf. Theory* **1** 416–23

[82] Pepper A, Tischler N and Pryde G J 2019 Experimental realization of a quantum autoencoder: the compression of qutrits via machine learning *Phys. Rev. Lett.* **122** 060501

[83] Romero J, Olson J P and Aspuru-Guzik A 2017 Quantum autoencoders for efficient compression of quantum data *Quantum Sci. Technol.* **2** 045001

[84] Tang E 2021 Quantum principal component analysis only achieves an exponential speedup because of its state preparation assumptions *Phys. Rev. Lett.* **127** 060503

[85] Miguel Arrazola J, Delgado A, Roy Bardhan B and Lloyd S 2020 Quantum-inspired algorithms in practice *Quantum* **4** 307

[86] Cotler J, Huang H-Y and McClean J R 2021 Revisiting dequantization and quantum advantage in learning tasks (arXiv:2112.00811)

[87] Huang H-Y *et al* 2022 Quantum advantage in learning from experiments *Science* **376** 1182–6

[88] Gordon M H, Cerezo M, Cincio L and Coles P J 2022 Covariance matrix preparation for quantum principal component analysis *PRX Quantum* **3** 030334

[89] Caro M C, Huang H-Y, Ezzell N, Gibbs J, Sornborger A T, Cincio L, Coles P J and Holmes Z 2022 Out-of-distribution generalization for learning quantum dynamics (arXiv:2204.10268)

[90] Agarwal R, Rethinasamy S, Sharma K and Wilde M M 2021 Estimating distinguishability measures on quantum computers (arXiv:2108.08406)

[91] Brun T A 2004 Measuring polynomial functions of states (arXiv:quant-ph/0401067)

[92] Tucci R R 2005 An introduction to Cartan's KAKD decomposition for QC Programmers (arXiv:quant-ph/0507171)

[93] Shende V V, Markov I L and Bullock S S 2004 Minimal universal two-qubit controlled-not-based circuits *Phys. Rev. A* **69** 062321

[94] Vatan F and Williams C 2004 Optimal quantum circuits for general two-qubit gates *Phys. Rev. A* **69** 032315

[95] Holmes Z, Arrasmith A, Yan B, Coles P J, Albrecht A and Sornborger A T 2021 Barren plateaus preclude learning scramblers *Phys. Rev. Lett.* **126** 190501

[96] Verstraete F, Cirac J I and Latorre J I 2009 Quantum circuits for strongly correlated quantum systems *Phys. Rev. A* **79** 032316

[97] Gibbs J, Holmes Z, Caro M C, Ezzell N, Huang H-Y, Cincio L, Sornborger A T and Coles P J 2022 Dynamical simulation via quantum machine learning with provable generalization (arXiv:2204.10269)

[98] Liu X, Liu G, Huang J and Wang X 2022 Mitigating barren plateaus of variational quantum eigensolvers (arXiv:2205.13539)

[99] Wilde M M 2011 From classical to quantum shannon theory (arXiv:1106.1445)

[100] Bures D 1969 An extension of Kakutani's theorem on infinite product measures to the tensor product of semifinite w^* -algebras *Trans. Am. Math. Soc.* **135** 199–212

[101] Bengtsson I and Zyczkowski K 2006 *Geometry of Quantum States: An Introduction to Quantum Entanglement* (Cambridge: Cambridge University Press)

[102] Hayashi M 2004 *Quantum Information Theory: Mathematical Foundation* 2nd edn (Berlin: Springer)

[103] Petz D and Sudár C 1996 Geometries of quantum states *J. Math. Phys.* **37** 2662–73

[104] Ginibre J 1965 Statistical ensembles of complex, quaternion and real matrices *J. Math. Phys.* **6** 440–9