


Universal Compiling and (No-)Free-Lunch Theorems for Continuous-Variable Quantum Learning

Tyler Volkoff^{1,*}, Zoë Holmes,² and Andrew Sornborger^{2,3}

¹*Theoretical Division, Los Alamos National Laboratory, Los Alamos, New Mexico, USA*

²*Information Sciences, Los Alamos National Laboratory, Los Alamos, New Mexico, USA*

³*Quantum Science Center, Oak Ridge, Tennessee 37931, USA*

 (Received 16 May 2021; revised 26 August 2021; accepted 29 September 2021; published 9 November 2021)

Quantum compiling, where a parameterized quantum circuit is trained to learn a target unitary, is an important primitive for quantum computing that can be used as a subroutine to obtain optimal circuits or as a tomographic tool to study the dynamics of an experimental system. While much attention has been paid to quantum compiling on discrete-variable hardware, less has been paid to compiling in the continuous-variable paradigm. Here we motivate several, closely related, short-depth continuous-variable algorithms for quantum compilation. We analyze the trainability of our proposed cost functions and numerically demonstrate our algorithms by learning arbitrary Gaussian operations and Kerr nonlinearities. We further make connections between this framework and quantum learning theory in the continuous-variable setting by deriving no-free-lunch theorems. These generalization bounds demonstrate a linear resource reduction for learning Gaussian unitaries using entangled coherent-Fock states and an exponential resource reduction for learning arbitrary unitaries using two-mode-squeezed states.

DOI: [10.1103/PRXQuantum.2.040327](https://doi.org/10.1103/PRXQuantum.2.040327)

I. INTRODUCTION

Progress in experimental implementations of quantum optical neural networks [1–3] and extensions of quantum machine learning frameworks to the continuous-variable (CV) setting [4–6] indicate that quantum photonics is a viable platform for near-term quantum algorithms. Variational quantum algorithms, where a problem-specific cost function is evaluated on a quantum computer [7,8], while a classical optimizer trains a parameterized quantum circuit to minimize this cost, have been implemented in photonic systems. For instance, the variational quantum eigensolver [9] and variational quantum unsampling [10], i.e., partial characterization of a unitary operator, have both been implemented on integrated photonic processors. Beyond the fundamental physical advantages of photonic systems, such as a well-characterized set of loss channels and the possibility of room-temperature operation, there are computational advantages to CV implementations of variational quantum algorithms such as the existence of efficient quantum error mitigation schemes [11–13].

An important computational task that CV quantum processors are well suited to is the variational compilation [14–16] of CV unitaries. The task is to optimize a parameterized quantum circuit to learn a given target unitary. The target unitary could take the form of a known gate sequence that one seeks to compile into a shorter depth, or more noise resistant, circuit. Hence, quantum compiling could be used as a subroutine to reduce the resources required to implement large-scale quantum algorithms. Alternatively, the target unitary could be the unknown dynamics of a quantum system. In this case, quantum compilation plays a role analogous to, but potentially less resource intensive than, a quantum sensing protocol [17] or unitary process tomography [18,19]. Specifically, our CV compiling algorithms make use of Gaussian measurements and CV resources such as intensity and quadrature squeezing, and so do not require preparation of exotic optimal probe states as in an optimal quantum sensing protocol, nor a large number of measured observables as in process tomography. In this sense, CV compiling provides a new tool for experimental physics.

In this paper, we establish frameworks for CV variational quantum compiling that are valid for arbitrary CV target unitaries. In contrast to the variational compiling method explored in Ref. [15], we include entanglement-enhanced methods that can be used to learn an entire unitary rather than just its action on a low lying subspace. We illustrate the wide applicability of our cost functions

*volkoff@lanl.gov

Published by the American Physical Society under the terms of the [Creative Commons Attribution 4.0 International](https://creativecommons.org/licenses/by/4.0/) license. Further distribution of this work must maintain attribution to the author(s) and the published article's title, journal citation, and DOI.

for CV quantum compiling by numerically demonstrating efficient learning of arbitrary single-mode Gaussian unitaries, the generalized beam splitter operation, and Kerr nonlinearities.

We further make connections between this framework and CV quantum learning theory by deriving “(no)-free-lunch” theorems. These analytic theorems specify the minimal training data required to learn CV unitary operators in increasingly general settings, providing fundamental bounds on the limits of quantum learning. In particular, the bounds highlight how utilizing entangled training states can reduce the amount of training data required to learn an unknown unitary and thus entanglement could be seen to provide a “free lunch.” We further use these results as an alternative motivation for the cost functions we propose for quantum compiling.

This manuscript is structured as follows. Section II provides a background to quantum compiling, including a discussion of its possible uses and a summary of previously proposed methods for discrete-variable quantum compilation. Section III presents our main results, including the cost functions we propose for CV quantum compiling and an analysis of their trainability. Section IV contains numerical implementations of our proposed CV learning algorithm. Section V presents our no-free-lunch theorems for CV learning. Section VI summarizes and discusses our results.

II. BACKGROUND

A. Applications of CV quantum compiling

The goal of CV quantum compiling is to take a (possibly unknown) unitary U and return a gate sequence V , executable on a CV quantum computer, that has approximately the same action as U on any given input state (up to possibly a global phase factor). Here we describe three possible applications of this subroutine.

1. Optimal circuit design

Quantum compilation could be used to variationally compile CV gate sequences to form optimal subcircuits. By optimal, we primarily mean short depth. However, compilation might also be used to find circuits that naturally compensate for systematic gate errors or that are more generally resistant to noise. The construction of such optimal circuits may prove critical for the successful implementation of larger scale algorithms, including proposals for generating optimal bosonic states in protocols for quantum metrology [20] and entanglement extraction [21].

2. Experimental quantum physics

More generally, variational quantum compilation could be used to learn the unknown unitary dynamics of a physical system. In the context of an optical system, one




FIG. 1. Learning experimental CV quantum systems. Here we sketch an experimental circuit to learn the unitary U implemented by a novel optical material (shown in orange) by training a parameterized quantum circuit $V^\dagger(\theta)$ implemented on an optical quantum computer (shown in gray). For specific details on the proposed circuit, and the cost it calculates, see Fig. 2(d).

might be interested in studying the optical properties of a new material as sketched in Fig. 1. For example, as discussed further in Sec. IV, one might use quantum compiling to estimate the Kerr effect in nonlinear optical media that cannot itself be directly implemented in a CV quantum circuit. In this manner, variational quantum compilation provides a new tool for experimental physics.

3. Structured learning

In discrete-variable systems, variational quantum compilation has proven useful for learning the spectral decomposition of a unitary operation. This in turn opens up the possibility of simulating beyond the coherence time of a quantum processor [22–24]. Similarly one could use discrete-variable quantum compiling to learn block decompositions of a given unitary that is useful to study the entanglement properties of a system. It would be interesting to explore whether CV quantum compiling could similarly be used to study the spectral or entanglement properties of a given CV unitary, or for simulating the dynamics of CV quantum systems [25–27].

B. Discrete-variable quantum compilation

Before presenting our algorithm for continuous-variable quantum compilation, let us first review the discrete-variable quantum-assisted quantum compilation (QAQC) algorithm of Ref. [14]. In QAQC a compilation is found by variationally searching for a gate sequence that minimizes the Hilbert-Schmidt test (HST) cost. This cost, which quantifies how close the compilation is to exact, can be written as the normalized Hilbert-Schmidt inner product between the target unitary U and possible compilation V ,

$$C_{\text{HST}}(V, U) := 1 - \frac{1}{d^2} |\text{Tr}(V^\dagger U)|^2. \quad (1)$$

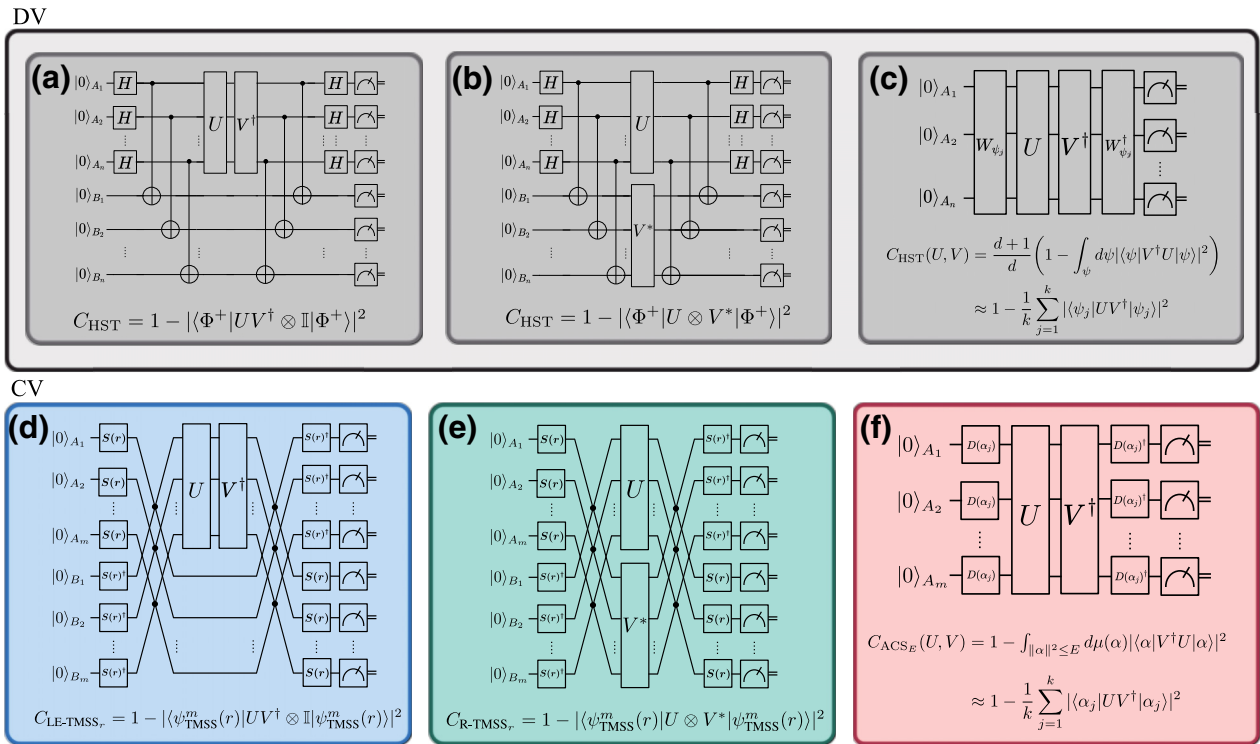


FIG. 2. Schematic of cost functions. In this figure we show how the three different cost functions we propose for CV quantum compiling (d)–(f) are related to three different ways of measuring the Hilbert-Schmidt test cost for DV compilation (a)–(c). The probability of measuring the all-zero state on all $2n$ qubits in (a) and (b) is equal to $\text{Tr}[UV^\dagger]/d^2$ and hence can be used to compute C_{HST} . In (c) W_{ψ_j} is the unitary that prepares the state $|\psi_j\rangle$, i.e., $W_{\psi_j}|0\rangle = |\psi_j\rangle$; therefore, the probability of measuring the all-zero state on n qubits is equal to $|\langle \psi_j | UV^\dagger | \psi_j \rangle|^2$. We could theoretically estimate C_{HST} by running this circuit over a Haar random ensemble of training states, but for large problems, this is exponentially inefficient. In (d) and (e) $S(r)$ is the unitary single-mode squeeze operator and \mathbb{X} is a 50:50 beam splitter that entangles the squeezed registers. The probability of measuring the all-zero state on all $2m$ modes in (d) and (e) is equal to $1 - C_{\text{LE-TMSS}}$ and $1 - C_{\text{R-TMSS}}$, respectively. In (f) $D(\alpha)$ is the unitary single-mode displacement operator. The mean probability of obtaining the all-zero state on m modes is equal to the average of the k values $|\langle \alpha_j | UV^\dagger | \alpha_j \rangle|^2$, which is used to estimate C_{ACS} .

This cost is faithful, vanishing if and only if U and V differ by a global phase factor, i.e., $V = e^{i\varphi} U$ for some $\varphi \in \mathbb{R}$. Therefore, by minimizing C_{HST} , we learn a unitary U that implements a target V up to a global phase.

The Hilbert-Schmidt test cost may be computed by the two closely related circuits shown in Figs. 2(a) and 2(b). To see how, we first note that

$$\frac{1}{d} \text{Tr}(V^\dagger U) = \langle \Phi^+ | V^\dagger U \otimes \mathbb{I} | \Phi^+ \rangle, \quad (2)$$

where $|\Phi\rangle_{\text{AB}}$ is the Bell entangled state of two qubit registers A and B of $n = \log_2 d$ qubits, i.e., $|\Phi\rangle_{\text{AB}} := \otimes_{j=1}^n |\Phi^+\rangle_{A_j B_j}$ with $|\Phi^+\rangle := (1/\sqrt{2})(|00\rangle + |11\rangle)$. It thus follows that we can write

$$C_{\text{HST}}(V, U) = 1 - |\langle \Phi^+ | V^\dagger U \otimes \mathbb{I} | \Phi^+ \rangle|^2 \quad (3)$$

and C_{HST} can be computed using the circuit shown in Fig. 2(a). Because of the ricochet property of the state $|\Phi\rangle$,

viz., $X \otimes \mathbb{I} | \Phi^+ \rangle = \mathbb{I} \otimes X^T | \Phi^+ \rangle$ for linear operator X , the Hilbert-Schmidt test cost can alternately be written as

$$C_{\text{HST}}(V, U) = 1 - |\langle \Phi^+ | U \otimes V^* | \Phi^+ \rangle|^2. \quad (4)$$

Thus, $C_{\text{HST}}(V, U)$ can also be computed with the target and ansatz unitaries applied in parallel, instead of in series, reducing the total circuit depth as shown in Fig. 2(b).

Finally, we note that the Hilbert-Schmidt test cost can be related to the average gate fidelity between U and V . Specifically, it can be shown [28,29] that

$$C_{\text{HST}}(U, V) = \frac{d+1}{d} [1 - \bar{F}(U, V)], \quad (5)$$

where

$$\bar{F}(U, V) := \int_{\psi} |\langle \psi | V^\dagger U | \psi \rangle|^2 d\psi \quad (6)$$

is the average fidelity of states acted upon by V versus those acted upon by U , with the average being over all pure

states according to the Haar measure. In theory, Eq. (5) provides a third way of measuring C_{HST} . One could perform a Loschmidt echo test, as shown in Fig. 2(c), using different input states that are sampled according to the Haar measure. However, in practice, this is not a viable training technique, since, as we discuss in Sec. V B, in order to fully learn U , the average would need to be taken over an exponentially large number of training states. Instead, the significance of Eq. (5) lies in the fact that it imbues C_{HST} with operational meaning for nonzero values since it entails that low cost values correspond to high average gate fidelities.

III. UNIVERSAL CONTINUOUS-VARIABLE QUANTUM COMPILING

A. Cost functions

For continuous-variable quantum compiling, we suppose that the target and compiled unitaries U and V act on m CV modes. More concretely, the mathematical setting for CV quantum compiling is the Hilbert space \mathcal{H} of m quantum harmonic oscillators, and the operator algebra $B(\mathcal{H})$ of bounded linear operators. In this infinite-dimensional space, the Hilbert-Schmidt inner product used in Eq. (1) is not well defined and hence cannot be used for CV quantum compilation. However, as indicated in Fig. 2, we can use generalizations of the three different ways in which Eq. (1) can be implemented, namely via Eqs. (3), (4), and (5), to define costs for CV quantum compiling. In contrast to the discrete-variable case where Eqs. (3), (4), and (5) are three equivalent ways to estimate the same cost, here the three costs Eqs. (10), (13), and (16) are fundamentally distinct. Thus, these costs do not have a direct operational meaning in terms of state-independent distinguishability of U and V . Correspondingly, the variational compiling algorithms defined by these costs can be viewed as heuristic optimization methods that are informed by the task of efficiently learning CV unitary operations. The subsections below introduce the cost functions for CV quantum compiling, and the connection between these functions and the task of unitary learning is detailed in Sec. V.

1. Loschmidt echo two-mode-squeezed state cost

Let us start by defining a CV generalization of Eq. (3). To do so, we first note that two-mode-squeezed states (TMSSs) are a natural analogue of Bell states for CV systems. The two-mode-squeezed state, acting between two m -mode registers A and B , is defined as

$$|\psi_{\text{TMSS}}^m(r)\rangle := \bigotimes_{j=1}^m |\psi_{\text{TMSS}}(r)\rangle_{A_j B_j} \quad (7)$$

$$\text{with } |\psi_{\text{TMSS}}(r)\rangle_{A_j B_j} \propto \sum_{n=0}^{\infty} (\tanh r)^n |n\rangle_{A_j} \otimes |n\rangle_{B_j},$$

where $\{|n\rangle\}_{n=0}^{\infty}$ is the Fock basis and r is a squeezing parameter. To highlight the connection between TMSSs and Bell states, it is helpful to consider its truncated variant

$$|\psi_{\text{TMSS}}^{\tau}(r)\rangle := \sqrt{\frac{1 - \tanh^{2\tau} r}{1 - \tanh^2 r}} \sum_{n=0}^{\tau-1} (\tanh r)^n |n\rangle \otimes |n\rangle, \quad (8)$$

which tends to the standard TMSS in the limit that τ tends to infinity, i.e., $\lim_{\tau \rightarrow \infty} |\psi_{\text{TMSS}}^{\tau}(r)\rangle = |\psi_{\text{TMSS}}(r)\rangle$. For finite τ , the truncated TMSS tends to a Bell state as r tends to infinity, that is,

$$\lim_{r \rightarrow \infty} |\psi_{\text{TMSS}}^{\tau}(r)\rangle = \frac{1}{\sqrt{\tau}} \sum_{n=0}^{\tau-1} |n\rangle \otimes |n\rangle = |\Phi^+\rangle. \quad (9)$$

In this sense, the TMSS may be viewed as a CV generalization of the Bell state.

More generally, TMSSs are highly entangled states that, by reducing the number of measurements necessary to attain a given signal-to-noise ratio, have proven to be an important resource in quantum metrology [30–32]. Moreover, TMSSs were numerically shown to be nearly optimal for measuring the fidelity of noisy CV quantum teleportation channels [33]. These examples suggest that TMSSs may also be valuable for the unitary channel discrimination task we consider here. This is confirmed in Sec. V, where we use the entanglement-enhanced no-free-lunch theorem of Ref. [34] to argue that training on a single TMSS minimizes the generalization error for learning a CV unitary.

This motivates our first proposed cost function to train an m -mode hypothesis unitary V to match an m -mode target unitary U as follows:

$$C_{\text{LE-TMSS},r}(V, U) := 1 - |\langle \psi_{\text{TMSS}}^m(r) | UV^\dagger \otimes \mathbb{1} | \psi_{\text{TMSS}}^m(r) \rangle|^2. \quad (10)$$

This is the CV analogue of Eq. (3) obtained by using an m -mode TMSS instead of a Bell state. We call this cost, which is evidently faithful by construction, the Loschmidt echo two-mode squeezed state (LE-TMSS) cost since it measures the inner product between $V \otimes \mathbb{1} | \psi_{\text{TMSS}}^m(r) \rangle$ and $U \otimes \mathbb{1} | \psi_{\text{TMSS}}^m(r) \rangle$ using the Loschmidt echo circuit sketched in Fig. 2(d).

To understand the structure of the circuit that we propose to measure $C_{\text{LE-TMSS},r}$, it is helpful to recall that the TMSS can be written as

$$|\psi_{\text{TMSS}}(r)\rangle_{A_j B_j} = e^{\pi(a_j b_j^\dagger - a_j^\dagger b_j)/4} S(-r) |0\rangle_{A_j} \otimes S(r) |0\rangle_{B_j}, \quad (11)$$

where a_j and b_j are the annihilation operators on the A_j and B_j modes, respectively, and $S(r)$ is the single-mode

squeezing operator [35]. It follows that an m -mode TMSS can be prepared across two m -mode registers A and B by first negatively squeezing the modes on register A and positively squeezing the modes on register B and then pairwise entangling the modes A_j and B_j (for $j = 1$ to $j = m$) using a network of 50:50 beam splitters.

As shown in Fig. 2(d), preparing a TMSS in this manner is the first step of the circuit to measure $C_{\text{LE-TMSS},r}$. The second step is to apply the target unitary U and the inverse of the ansatz V^\dagger to register A . The final step is to implement the inverse of the m -mode TMSS preparation in order to measure the overlap with the m -mode TMSS. This is done by first inverting the beam splitter network and then reversing the initial local squeezing. The inverse squeezing can be carried out either by a two-step process consisting of an active optical unitary followed by an on-off photodetection measurement, or in a one-step process by an ideal general-dyne measurement [36]. The probability of obtaining the measurement outcome in which all $2m$ modes are in the $|0\rangle$ state, i.e., the vacuum state, is equal to $|\langle \psi_{\text{TMSS}}^m(r) | UV^\dagger \otimes \mathbb{1} | \psi_{\text{TMSS}}^m(r) \rangle|^2$. Hence, this circuit can be used to measure $C_{\text{LE-TMSS},r}$, as claimed.

2. Ricocheted two-mode-squeezed state cost

Unlike the Bell states utilized in discrete-variable quantum compiling algorithms, the TMSS only satisfies an approximate ricochet property for finite r . That is, with $|\psi_{\text{TMSS}}^m(r)\rangle$ defined as in Eq. (7),

$$UV^\dagger \otimes \mathbb{1}_B |\psi_{\text{TMSS}}^m(r)\rangle \approx U_A \otimes V_B^* |\psi_{\text{TMSS}}^m(r)\rangle \quad (12)$$

with the exact property only holding in the limit that $r \rightarrow \infty$ or for specially chosen U and V . Consequently, the Ricocheted version of the two-mode-squeezed state cost function, i.e.,

$$C_{\text{R-TMSS},r}(V, U) := 1 - |\langle \psi_{\text{TMSS}}^m(r) | U_A \otimes V_B^* | \psi_{\text{TMSS}}^m(r) \rangle|^2, \quad (13)$$

is equal to Eq. (10) in the limit $r \rightarrow \infty$. The circuit for computing Eq. (13), which is shown in Fig. 2(e), is identical to the circuit used to measure $C_{\text{LE-TMSS}}$ but with the target and ansatz unitaries prepared in parallel rather than series. The difference between cost functions (10) and (13) depends on V . In Appendix A we show how the cost functions differ in expectation over finite rank V . The calculation shows that even if the size of V increases multiplicatively, it is sufficient to increase the squeezing parameter r additively in order to make the cost functions (10) and (13) approximately equal.

Although the $C_{\text{R-TMSS},r}$ cost can be computed by a simple circuit, it has a drawback that its minimum need not be zero when optimizing V for a given U . Hence, when used for variational compiling, it will be hard to determine when to

terminate the optimization loop. It is therefore helpful to define a normalized version of Eq. (13):

$$\tilde{C}_{\text{R-TMSS},r}(V, U) := 1 - \frac{|\langle \psi_{\text{TMSS}}^m(r) | U_A \otimes V_B^* | \psi_{\text{TMSS}}^m(r) \rangle|^2}{\mathcal{N}_U \mathcal{N}_V}. \quad (14)$$

Here the normalization terms,

$$\mathcal{N}_X := |\langle \psi_{\text{TMSS}}^m(r) | X_A \otimes X_B^* | \psi_{\text{TMSS}}^m(r) \rangle| \quad (15)$$

for $X = U$ and $X = V$, can be calculated using the same circuit to measure $C_{\text{R-TMSS},r}$. As shown in Appendix B, this normalized cost $\tilde{C}_{\text{R-TMSS},r}$ is faithful, vanishing if and only if U and V agree up to a global phase.

Given the need to evaluate the normalization terms, as well as the original cost term, this cost is slightly more resource intensive than the LE-TMSS cost. However, the reduction in circuit depth achieved by using the approximate ricochet property may compensate for this in experimental contexts where coherence lifetimes are short.

3. Averaged coherent states cost

It is not possible to define a cost that is directly analogous to Eq. (5) in a CV context as there is no direct equivalent to the Haar measure for CV states because of the infinite dimensionality of the Hilbert space for CV systems. Instead, one can consider averaging over a family of states up to a specific energy bound. In Ref. [15] a cost is defined in this manner as an average over Fock states. However, large Fock states are hard to produce experimentally, and so we argue that a more natural choice, given the ease with which they can typically be produced in the laboratory, is coherent states. With this in mind, one could consider using the cost function

$$C_{\text{ACS}_E}(V, U) := 1 - \int_{\|\alpha\|^2 \leq E} d\mu(\alpha) |\langle \alpha | V^\dagger U | \alpha \rangle|^2, \quad (16)$$

where $d\mu(\alpha)$ is a normalized measure on the set of m -mode coherent states $|\alpha\rangle$ with energy [37] less than E . Each of the coherent state overlaps in Eq. (16) can be computed using local heterodyne measurements on the m modes. That Eq. (16) is faithful can be seen from the fact that if it takes the value 0, the modulus of the Q symbol of the unitary operator $V^\dagger U$ is equal to 1 almost everywhere on the domain, from which it follows that $V^\dagger U = e^{i\phi}$ due to the overcompleteness of coherent states [38].

In practice, this cost, which we call the averaged coherent state (ACS) cost, can only be estimated by sampling k

coherent states with energy less than E , i.e., using

$$C_{ACS_E}(V, U) \approx 1 - \frac{1}{k} \sum_{j=1}^k |\langle \alpha_j | V^\dagger U | \alpha_j \rangle|^2, \quad (17)$$

where $\|\alpha_j\|^2 < E$ for all j . In Sec. V we use a NFL theorem for Gaussian operations to argue that $k = 2m$ training states will suffice to learn any Gaussian unitary U . In Sec. IV we provide numerics that suggest that to learn weakly non-Gaussian operations, in particular a small Kerr nonlinearity, $k = 2m$ modes is also sufficient. However, in general, to learn an arbitrary operation, we expect that k will scale with E .

Thus, in general, estimating C_{ACS_E} will be more resource intensive than the TMSS costs, $C_{LE-TMSS_r}$ and C_{R-TMSS_r} , in the sense that it requires a larger number of cost evaluations. However, C_{ACS_E} does not require generating large, highly entangled two-mode-squeezed states, and therefore may in some contexts be less experimentally demanding. In particular, we expect this cost to be most useful for learning (approximately) Gaussian operations where the number of training states required is reduced.

B. Trainability

For a variational quantum algorithm to run successfully, i.e., for it to be possible to minimize cost and thereby find the optimum solution, the cost landscape must have sufficiently large gradients to allow for training. Recently, it has been shown that discrete-variable variational quantum algorithms can exhibit so-called ‘‘barren plateaus,’’ where, under certain conditions, the gradient of the cost function vanishes exponentially with the size of the system [39–51]. Preliminary results further indicate continuous-variable systems [52] may exhibit an analogous barren plateau phenomenon where the cost gradients vanish exponentially with the number of system modes. On such barren plateau landscapes (potentially untenably) precise measurements are required to determine the direction of steepest descent and navigate to the minimum. Thus, for any learning algorithm to be *scalable* to large problem sizes, it is essential to use a cost that does not exhibit a barren plateau.

Even from basic examples, one can see that the cost functions for CV quantum compiling exhibit a barren plateau. To demonstrate this, we focus on the Loschmidt echo TMSS cost, but analogous arguments follow for the ricocheted TMSS cost and the averaged coherent state cost. Consider using the Loschmidt echo TMSS cost to compile the m -mode identity operation using the ansatz composed of a product of phase gates, i.e., $V = e^{i \sum_{j=1}^m \phi_j a_j^\dagger a_j}$, where the ϕ_j are uniform in $[-\pi, \pi]$. Then the cost takes

the form

$$\begin{aligned} C_{LE-TMSS_r}(\phi) &= 1 - \frac{1}{\cosh^{4m} r} \prod_{j=1}^m [(1 - 2 \cos \phi_j \tanh^2 r + \tanh^4 r)^{-1}]. \end{aligned} \quad (18)$$

It follows (examining ϕ_1 without loss of generality) that

$$\begin{aligned} E(|\partial_{\phi_1} C_{LE-TMSS_r}|) &= \frac{1}{(2\pi)^m} \int_0^{2\pi} d\phi |\partial_{\phi_1} C_{LE-TMSS_r}| \\ &= \left(\frac{2}{\pi(1 + 2 \sinh^2 r)^2} \right)^m \frac{\tanh^2 r}{1 + \tanh^4 r}, \end{aligned} \quad (19)$$

which vanishes exponentially with the number of modes m . It therefore follows from Chebyshev’s inequality

$$P(|\partial_{\phi} C_{LE-TMSS_r}| > \epsilon) \leq \frac{E(|\partial_{\phi} C_{LE-TMSS_r}|)}{\epsilon} \quad (20)$$

that the probability that the cost gradient deviates from zero vanishes exponentially. Thus, the landscape exhibits a barren plateau [52]. If r is allowed to vary with m then taking sublinear scaling of the total squeezing, e.g., local squeezing $r(m) = O(\ln m^\alpha / m)$, $\alpha > 0$, causes Eq. (20) to vanish only polynomially.

For fixed r , the barren plateau phenomenon can be circumvented by using a local variant of our proposed costs. Analogously to the local version of the Hilbert-Schmidt test introduced in Ref. [14], where pairs of qubits, rather than all $2n$ qubits, are measured to compute the local cost, our proposed local TMSS costs can be calculated from measurements on pairs of CV modes instead of $2m$ CV modes. Specifically, as shown in Fig. 3(a), the local version of the Loschmidt echo TMSS cost is defined as

$$C_{LE-TMSS_r}^{(L)}(V, U) := 1 - \frac{1}{m} \sum_{j=1}^m \Pr(00)_{A_j B_j}, \quad (21)$$

where $\Pr(00)_{A_j B_j}$ is the probability of observing outcome 00 on the pair of modes $A_j B_j$ from registers A and B in Fig. 2(d). This cost function can be shown to be faithful using the same probability theoretic argument used to prove the faithfulness of the local Hilbert-Schmidt test cost in Ref. [14].

The local cost $C_{LE-TMSS_r}^{(L)}$ can be expressed as a sum of entanglement fidelities. To see how, first note that the local

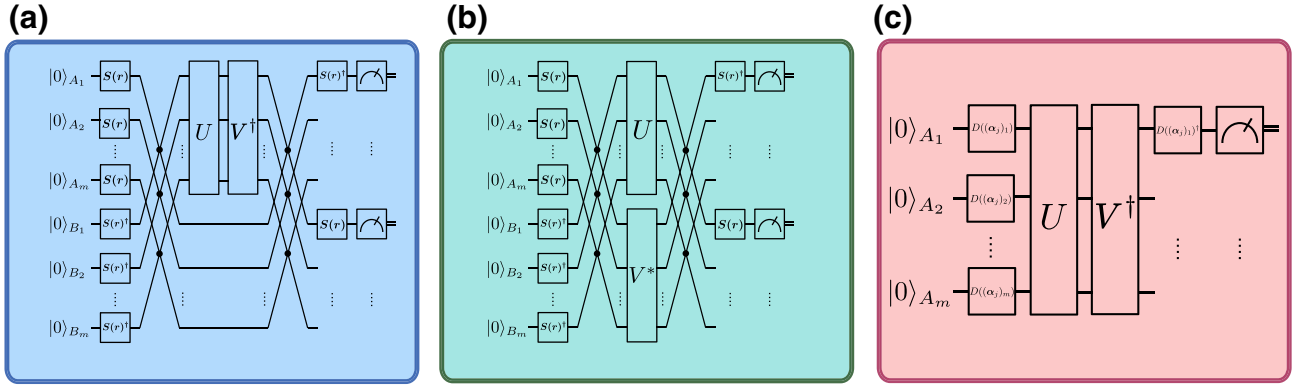


FIG. 3. Local cost functions. Here we show the circuit diagrams for the local versions of the (a) Loschmidt echo TMSS cost $C_{\text{LE-TMSS}}^{(L)}$, (b) ricocheted TMSS cost $C_{\text{R-TMSS}}^{(L)}$, and (c) averaged coherent state cost $C_{\text{ACS}_E}^{(L)}$. Crucially, in contrast to their respective global variants, only a pair of modes (in the case of $C_{\text{LE-TMSS}}^{(L)}$ and $C_{\text{R-TMSS}}^{(L)}$) and a single mode (in the case of $C_{\text{ACS}_E}^{(L)}$) is measured per circuit evaluation.

marginal states of $UV^\dagger \otimes \mathbb{I}_B |\psi_{\text{TMSS}}^m(r)\rangle$ on $A_j B_j$, i.e.,

$$\rho_{A_j B_j} := \text{Tr}_{\overline{A_j} \overline{B_j}} [(\mathcal{E}_{UV^\dagger} \otimes \mathbb{I}_B) |\psi_{\text{TMSS}}^m(r)\rangle \langle \psi_{\text{TMSS}}^m(r)|],$$

where $\overline{A_j}$ ($\overline{B_j}$) is the complement of A_j (B_j) in mode set A (B) and $\mathcal{E}_{UV^\dagger}(\cdot) = UV^\dagger(\cdot)VU^\dagger$, can be written as

$$\rho_{A_j B_j} = \text{Tr}_{A_j} [\mathcal{E}_{UV^\dagger} \otimes \mathbb{I}_{B_j} \{\rho_{A_j B_j}^{\text{TMSS}}(r) \otimes \rho_{\beta(r)}^{\otimes m-1}\}]. \quad (22)$$

Here $\rho_{\beta(r)}$ is a thermal state at the inverse temperature $\beta(r) := -2 \ln \tanh r$,

$$\rho_{\beta(r)} := \frac{1}{\cosh^2 r} \sum_{n=0}^{\infty} \tanh^2 r |n\rangle \langle n|, \quad (23)$$

and $\rho_{A_j B_j}^{\text{TMSS}}(r) := |\psi_{\text{TMSS}}^m(r)\rangle \langle \psi_{\text{TMSS}}^m(r)|_{A_j B_j}$. It follows that $C_{\text{LE-TMSS}}^{(L)}$ can be written as

$$C_{\text{LE-TMSS}_r}^{(L)}(V, U) := 1 - \frac{1}{m} \sum_{j=1}^m F_j, \quad (24)$$

where F_j is the entanglement fidelity of the channel

$$\mathcal{E}_j(\rho_{A_j}) := \text{Tr}_{\overline{A_j}} UV^\dagger [\rho_{A_j} \otimes (\rho_{\beta(r)})^{\otimes m-1}] VU^\dagger \quad (25)$$

with respect to the TMSS. That is,

$$F_j := \text{Tr}[\rho_{A_j B_j}^{\text{TMSS}}(\mathcal{E}_j \otimes 1_{B_j})(\rho_{A_j B_j}^{\text{TMSS}})]. \quad (26)$$

Thus, not only is the local cost faithful, it also has a natural conceptual interpretation.

Crucially, the local TMSS cost (24) appears not to exhibit a barren plateau. For example, for the problem

of compiling the identity with multimode phase shifters considered at the beginning of this section, one obtains

$$E(|\partial_{\phi_1} C_{\text{LE-TMSS}_r}^{(L)}|) = \frac{2}{\pi m \cosh^4 r (1 + \tanh^2 r)^2} \times \left(m - 1 + \frac{\tanh^2 r}{1 + \tanh^4 r} \right), \quad (27)$$

which, for fixed r , is constant as $m \rightarrow \infty$.

We note that, for a fixed number of modes m , the costs $C_{\text{LE-TMSS}_r}$ and $C_{\text{LE-TMSS}_r}^{(L)}$ concentrate to 1 when the squeezing parameter r is large. It follows that the gradients of $C_{\text{LE-TMSS}_r}$ and $C_{\text{LE-TMSS}_r}^{(L)}$, as seen from Eqs. (19) and (27), vanish exponentially with r . Consequently, training becomes exponentially more resource intensive for larger r . A similar exponential vanishing with respect to r was observed for approximations of CV energy-constrained channel fidelities that compare the actions of CV channels on two-mode-squeezed states [33]. This vanishing gradient problem is conceptually different to the barren plateau phenomenon that may be resolved using a local cost. In Sec. IV, we propose a practical resolution for this vanishing gradient problem.

Finally, we note that, for the example of compiling the identity operation considered earlier, the averaged coherent state cost function C_{ACS_E} in Eq. (16), and its approximation in Eq. (17), do not exhibit barren plateaus if the energy bound E is taken to depend on the mode number m in such a way that the maximal energy per mode $E(m)/m$ grows sublinearly as a function of m (see Sec. 2 of Ref. [52]). However, we expect that more general compiling problems, such as Gaussian compiling or compiling of Kerr nonlinearities, will exhibit barren plateaus. Such trainability issues could again be mitigated by defining a local version of C_{ACS_E} . A natural choice in local cost would be [analogously to Eq. (21)] to compute a spatial average

of the probability of measuring the vacuum state on each of the modes at the end of the circuit in Fig. 2(f). More concretely, one could use

$$C_{\text{ACSE}}^{(L)}(V, U) := 1 - \frac{1}{k} \sum_{j=1}^k \text{Tr}[O_{\text{Local}}^{(j)} V^\dagger U |\alpha_j\rangle \langle \alpha_j| U^\dagger V],$$

$$O_{\text{Local}}^{(j)} = \frac{1}{m} \sum_{\ell=1}^m |(\alpha_j)_\ell\rangle \langle (\alpha_j)_\ell|_{A_\ell} \otimes \mathbb{I}_{\overline{A_\ell}}$$

subject to $\|\alpha_j\|^2 \leq E$ for all j ,

where $|\alpha_j\rangle = \bigotimes_{\ell=1}^m |(\alpha_j)_\ell\rangle_{A_\ell}$ is a coherent state in the $2m$ -dimensional phase space, and we have used the discrete version of $C_{\text{ACSE}}(V, U)$ in Eq. (17). Computation of one term in the double sum defining $C_{\text{ACSE}}^{(L)}$ is shown in Fig. 3(c).

IV. NUMERICAL IMPLEMENTATIONS

Here we present results for implementing CV quantum compilation to learn commonly encountered CV operations. In particular, we focused on learning arbitrary single-mode Gaussian operations, Kerr nonlinearities, and a general beam splitter operation. In each case, we performed continuous parameter optimization in order to minimize the TMSS cost function, Eq. (13). We focus on the Loschmidt echo variant of the cost, but similar results are obtained for the Ricocheted variant. We note that it is unnecessary to use the local version of the cost here since, for these proof-of-principle implementations, we consider learning single- and two-mode unitaries for which the cost gradients are expected to be manageable even with a global cost.

Given the close connections between the TMSS cost and the HST cost for large r , and the operational meaning of C_{HST} as a measure of the average fidelity between U and V , ideally we would use a large r value, i.e., large squeezing, to learn U . However, as discussed in Sec. III B, and as demonstrated numerically in Fig. 4, the landscape of the TMSS cost becomes overwhelmingly flat for large r , making it difficult to train. We therefore found it more effective to train initially using a small r value. Then once reasonably accurate pre-trained parameters have been obtained using a small r , we trained on a larger r to refine the quality of the solution.

A. Gaussian operations

An arbitrary single-mode Gaussian operation

$$U_{\text{Gaus}}(\alpha, \beta, \phi) := e^{-iH_{\text{Gaus}}(\alpha, \beta, \phi)} \quad (28)$$




FIG. 4. Cost landscapes. The cost landscape for learning (a) an arbitrary Gaussian operation and (b) a $\chi = 3$ Kerr nonlinearity using four layers of the general ansatz defined in Eq. (31). Here ϵ is a noise parameter that determines the deviation of the ansatz parameters, θ , from the optimum parameters, θ^{opt} . Specifically, we set $\theta_k = \theta_k^{\text{opt}} + \epsilon R$, where R is a random number between -1 and 1 .

is generated by the quadratic Hamiltonian

$$H_{\text{Gaus}}(\alpha, \beta, \phi) := \alpha a + \alpha^* a^\dagger + \beta a^2 + \beta^* a^{\dagger 2} + \phi a^\dagger a, \quad (29)$$

where α and β are arbitrary complex numbers and ϕ is an arbitrary real number. We generated a random target Gaussian operation $U_{\text{targ}} := U_{\text{Gaus}}(\alpha_{\text{targ}}, \beta_{\text{targ}}, \phi_{\text{targ}})$ by choosing $\text{Re}(\alpha_{\text{targ}})$, $\text{Im}(\alpha_{\text{targ}})$, $\text{Re}(\beta_{\text{targ}})$, and $\text{Im}(\beta_{\text{targ}})$ randomly in the range $[0, 1]$ and ϕ_{targ} in the range $[0, 2\pi]$. We then used $C_{\text{LE-TMSS}}$ to learn U_{targ} using an ansatz of the same form. That is, using an ansatz of the form $V_{\text{anz}} = U_{\text{Gaus}}(\alpha, \beta, \phi)$, where α , β , and ϕ are parameters to be variationally learnt. Since U_{Gaus} is readily factorizable into the product of displacement, squeezing, and phase operations, this ansatz can be straightforwardly implemented using standard gates on a CV quantum computer.

The results of learning [53] an arbitrary Gaussian operation are shown in the top row of Fig. 5. To quantify the quality of the optimization, we take the optimal parameters obtained at each iteration of the optimization algorithm and plot both the Hilbert-Schmidt cost, C_{HST} , and the errors in the individual optimized parameters. As shown in Fig. 5(a), we start with $r = 0.1$ and successfully optimize the TMSS cost (using the COBYLA algorithm) down to 10^{-6} . This corresponds to errors in the HST cost and individual parameters in the region of 10^{-1} to 10^{-4} . We then took the optimal parameters from minimizing the TMSS cost with $r = 0.1$ and optimized using the HST cost with

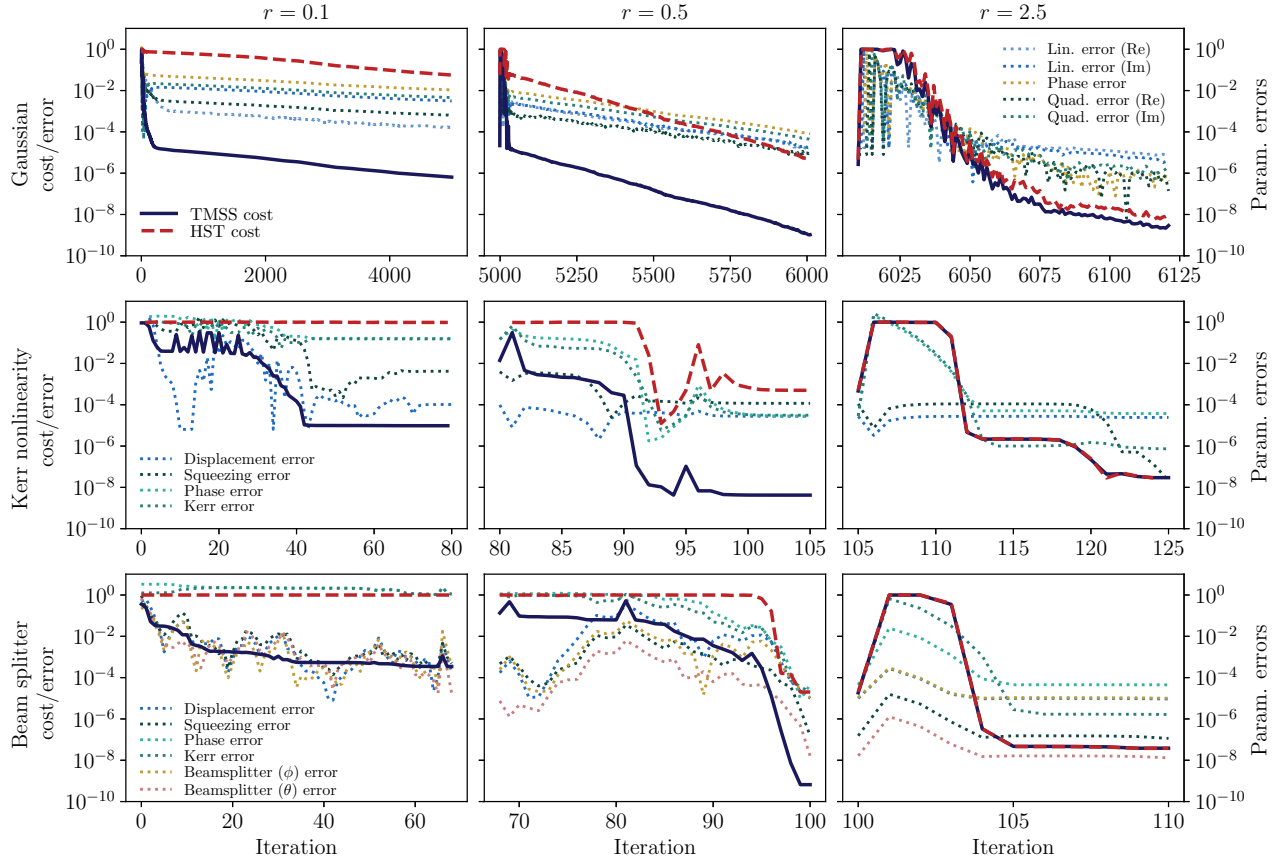


FIG. 5. Learning using TMSS cost. The $C_{\text{LE-TMSS}_r}$ cost (blue) as a function of iteration for learning a Gaussian (top row), Kerr nonlinearity (middle row), and beam splitter (bottom row). To mitigate the problem of vanishing cost gradients for large r , we take a perturbative approach starting with $r = 0.1$ (left column), and after convergence increasing r to 0.5 (middle column) and then 2.5 (right column). To quantify the quality of the optimization, we take the optimal parameters obtained at each iteration of the optimization algorithm and plot both the Hilbert-Schmidt cost C_{HST} (red dashed) and the errors in the individual optimized parameters (dotted). The parameter errors are given in natural units with $\hbar = 1$ and the mass and frequency of the modes equal to 1.

$r = 0.5$. The cost value and parameter errors initially go sharply up (because the old parameters that optimized the cost with $r = 0.1$ are no longer optimal) before decreasing again as the new cost is optimized. After optimizing with $r = 2.5$ we get both the TMSS and HST costs down to 10^{-9} with errors in the individual parameters in the region of 10^{-5} . Thus, Fig. 5 both demonstrates the effectiveness of our perturbative strategy and highlights how the difference between C_{TMSS} and C_{HST} decreases with increasing r .

B. Kerr nonlinearity

The second optimization task we consider is learning a Kerr nonlinearity of the form

$$U_{\text{Kerr}}(\chi) := e^{-i\chi(a^\dagger a)^2}. \quad (30)$$

Since there is no simple ansatz that can capture an arbitrary non-Gaussian operation, in this case we use the general layered ansatz advocated in Refs [54,55]. This ansatz is

composed of multiple layers that each consist of a displacement, squeeze, phase shift, and nonlinear Kerr shift. That is, a single layer is of the form

$$V_{\text{layer}}(\alpha, \beta, \phi, \chi) := U_{\text{Kerr}}(\chi)R(\phi)D(\alpha)S(\beta) \quad (31)$$

and the total ansatz is composed of a product of L such layers, $V_{\text{ans}}(\theta) := \prod_{l=1}^L V_{\text{layer}}(\alpha_l, \beta_l, \phi_l, \chi_l)$, where $\theta := \{\alpha_l, \beta_l, \phi_l, \chi_l\}_{l=1}^L$. Since the gates in every layer constitute a universal set [55], this ansatz can be used to implement any single-mode quantum operation.

We focus on the task of learning a large Kerr nonlinearity (of, perhaps, some new, yet to be classified, material). To make this task both nontrivial and physically pertinent, we suppose that the Kerr nonlinear components used as part of the ansatz are limited to implementing some maximum nonlinearity that is less than that of the target nonlinearity. Specifically, we suppose that the components of θ are bounded between 0 and 1, and we consider trying to learn $\chi_{\text{targ}} = 3$

using a four-layered ansatz. To perform the optimization, we employ the gradient-based limited-memory Broyden–Fletcher–Goldfarb–Shanno (BFGS) algorithm.

To assess the performance of the optimization in Fig. 5, we again plot the HST cost as well as a measure of the error in the individual parameters. Given the noncommutativity of the displacement, squeezing, phase shift, and Kerr operations, there are multiple possible choices in the parameters θ such that $V(\theta) = U_{\text{Kerr}}(\chi_{\text{targ}})$. Despite this, in practice, we find that the optimization algorithm found the “obvious” solution where the displacement parameters α_l , squeezing parameters β_l , and phase shift parameters ϕ_l each sum to zero and the Kerr nonlinearity parameters χ_l summed to χ_{targ} . We therefore take the difference between these values (i.e., $|\sum_{l=1} \alpha_l - 0|$, $|\sum_{l=1} \beta_l - 0|$, $|\sum_{l=1} \phi_l - 0|$, and $|\sum_{l=1} \chi_l - \chi_{\text{targ}}|$) as the measure of our displacement, squeezing, phase, and Kerr errors, respectively.

Similarly to the Gaussian case we find that starting with a small r allows for successful training. Then increasing the value of r improves the quality of the training in the sense that the HST error and parameter errors can be further decreased. We achieve a final TMSS and HST cost of 10^{-8} and parameter errors of approximately 10^{-5} .

C. Beam splitter operation

Finally, we attempt to learn a beam splitter operation of the form

$$U_{\text{BS}}(\theta, \phi) = e^{\theta(ab^\dagger e^{i\phi} - a^\dagger b e^{-i\phi})} \quad (32)$$

for a two-mode system with annihilation operators a and b , respectively, where θ and ϕ are randomly chosen phases in the range $[0, 2\pi]$. To learn this operation, we use a single-layer ansatz of the form

$$V_{\text{layer}}^{(12)}(\theta) := U_{\text{BS}}(\theta, \phi) V_{\text{layer}}^{(1)}(\alpha_1, \beta_1, \phi_1, \chi_1) \times V_{\text{layer}}^{(2)}(\alpha_2, \beta_2, \phi_2, \chi_2), \quad (33)$$

where $V_{\text{layer}}^{(1)}$ and $V_{\text{layer}}^{(2)}$ indicate the single-mode gate sequence defined in Eq. (31) on the first and second modes, respectively. The optimization was successful, with the TMSS and HST costs reduced to approximately 10^{-7} .

We note that, while it may superficially appear from Fig. 5 that fewer iterations are required to learn the Kerr nonlinearity and beam splitter operation than an arbitrary Gaussian operation, this is a feature of our choice in optimization algorithm. Namely, the BFGS algorithm uses a gradient-based approach that involves evaluating the cost n_{param} times, where n_{param} is the number of parameters that need to be learnt at every iteration step. Once this is accounted it requires more cost evaluations to learn the general beam splitter or a Kerr nonlinearity than to learn a

Gaussian operation. This is precisely as one would expect since these are more complex optimization problems.

V. NO-FREE-LUNCH THEOREMS FOR CV QUANTUM LEARNING

In classical machine learning, the no-free-lunch (NFL) theorems consider the task of learning a target function f , where f maps a discrete input set X to a discrete output set Y (both of size d). The learning is performed using a training set S consisting of $|S|$ input-output training pairs,

$$S = \{(x_j, f(x_j)) : x_j \in X\}_{j=1}^{|S|}. \quad (34)$$

In the limit of perfect learning, one assumes that it is possible to train a hypothesis function h_S to match the target function f on all training pairs in S . The no-free-lunch theorems then quantify the “generalization error,” i.e., how well the hypothesis function matches the target on unseen data. In general terms, the theorems demonstrate that the generalization error of a given learning algorithm is not less than that of a random learning algorithm in expectation over target functions f [56–60]. That is, the average performance of a learning algorithm is determined not by the choice in learning algorithm but rather by the amount of training data $|S|$.

Specifically, the generalization error can be quantified by the *risk* function

$$R_f(h_S) = \sum_{x \in X} \pi(x) \mathbf{1}[f(x) \neq h_S(x)], \quad (35)$$

where $\mathbf{1}[S]$ is the indicator function taking value 1 (0) if condition S is satisfied (not satisfied). This is the probability that the hypothesis function $h_S(x)$ and target function $f(x)$ differ across X , the domain of f , when x is sampled from the uniform probability distribution $\pi(x)$. The average risk, averaged over training sets S and functions f , can, for any optimization method, be lower bounded as [60]

$$E_f[E_S[R_f(h_S)]] \geq \left(1 - \frac{1}{d}\right) \left(1 - \frac{|S|}{d}\right). \quad (36)$$

Hence, the average risk is determined by the number of training pairs $|S|$, vanishing if and only if S spans the full domain of f , i.e., if $|S| = d$.

Similar NFL theorems exist for finite-dimensional quantum circuit learning, in which a target function f corresponds to a unitary quantum channel U and the training set is generalized to a set of quantum state pairs $S = \{|\psi_j\rangle, U|\psi_j\rangle\}_{j=1}^{|S|}$. By defining the generalization error using a suitable distance on quantum state space, it is shown that in general an exponential number of training states, $|S| \sim 2^n$, are required to learn an n qubit unitary [61].

TABLE I. No-free-lunch theorems for CV unitary learning.

Target unitary	Training set S	Entangled training	NFL	Cost motivated
Linear optical	Coherent states	No	$\frac{1}{2} - S /4m$	C_{ACS_E} (17)
Linear optical	Coherent-Fock states	Yes	$\frac{1}{2} - \tau S /4m$	C_{ECFS} (47)
Fock truncated	Schmidt rank τ TMSS	Yes	$1 - (\tau^2 S ^2 + d + 1)/d(d + 1)$	$C_{\text{LE-TMSS}}$ (10)

Furthermore, by allowing the training set to consist of pairs of states that are entangled with a reference system, i.e., $S = \{|\psi_j\rangle, U \otimes \mathbb{I}_{\mathcal{R}}|\psi_j\rangle\}_{j=1}^{|S|}$, where $|\psi_j\rangle \in \mathcal{H} \otimes \mathcal{H}_{\mathcal{R}}$ are entangled pure states of Schmidt rank τ , an entanglement-enhanced quantum no-free-lunch theorem can be derived. In this case, the lower bound of the expected error of a quantum learning algorithm, over all target unitaries U is reduced linearly in τ [34]. This has the important practical implication that, by using entanglement as a resource, the number of unique input-output state pairs needed to learn a target unitary, U , may be exponentially reduced in the limit of perfect learning.

In Sec. V A we derive NFL theorems in a restricted setting where the task is learning a linear optical unitary operation. Specifically, Theorems 1 and 2 quantify learning with classical training data (coherent states) and quantum training data (entangled coherent-Fock states), respectively. Section V B shows how the entanglement-assisted NFL theorem of Ref. [34] can be applied in an unrestricted CV learning setting. We further discuss how CV quantum NFL theorems can be used to motivate cost functions for CV quantum compiling. These results are summarized in Table I.

A. Learning linear optical unitaries from Gaussian training data

Linear optical unitaries capture the dynamics of multimode beam splitters, phase shifters, and displacement operators. Such unitaries, on m CV modes, can be written in the form $U = e^{iH}$, where $H = H^\dagger$ and $[H, \sum_{j=1}^m a_j^\dagger a_j] = 0$. Here we consider the task of training a hypothesis unitary V_S to emulate a target linear optical unitary U using a set of training data S composed of m -mode coherent states. We analyze the expected performance of a generic learning algorithm, over all target linear optical unitaries U and all training sets containing $|S|$ training states.

To fix the notation, an m -mode coherent state with mean vector w is written $|w\rangle$. Here w is a row vector in \mathbb{R}^{2m} given by $w = \langle R \rangle$ with $R = (q_1, p_1, \dots, q_m, p_m)$ the vector of canonical quadrature operators. The action of the target linear optical unitary U on $|w\rangle$ is given by $U|w\rangle = |wO\rangle$, where O is a $2m \times 2m$ orthogonal matrix. The set of $2m$ by $2m$ orthogonal matrices will be denoted $\text{Orth}(2m)$. Equipped with this notation, the training set to learn a linear optical unitary U using $|S|$ pairs of m -mode coherent

states can be written as

$$S = \{(w_j, w_j O)\}_{j=1}^{|S|} \in (\mathbb{R}^{2m} \times \mathbb{R}^{2m})^{\times |S|}. \quad (37)$$

Similarly, the action of the hypothesis linear optical unitary V_S on $|w\rangle$ can be written as $V_S|w\rangle = |wT_S\rangle$, where $T_S \in \text{Orth}(2m)$. We focus on the limit of perfect learning and assume that the learning algorithm outputs an orthogonal matrix T_S that agrees perfectly with O on all coherent states w_j in the training set. That is, we assume that $w_j T_S = w_j O$ for the training data mean vectors $(w_j, w_j O) \in S$.

To quantify how well the hypothesis unitary V_S matches the target unitary U on all possible coherent states, i.e., not just the training states, we define a risk function. To do so, we utilize a simple loss function of the form $L(y, z) = \|y - z\|^2$, where $y = xO$ and $z = xT_S$ are the output vectors of the target and hypothesis orthogonal matrices, respectively. Throughout this section, $\|\cdot\|$ refers to the 2-norm on the Euclidean space \mathbb{R}^{2m} . The total risk is then defined as the average loss over a multivariate Gaussian distribution of input vectors x , i.e., over the distribution $\pi(x) = e^{-\|x\|^2/2\sigma} / (2\pi\sigma)^m$. The total risk thus takes the form

$$R_O(T_S) = \frac{1}{8m\sigma} \int d^{2m}x \pi(x) L(xT_S, xO), \quad (38)$$

where the normalization factors have been chosen to ensure that $R_O(T_S)$ takes values between 0 and 1. In essence, $R_O(T_S)$ is a measure of how well T_S matches O isotropically in phase space. Risk values of $R_O(T_S) = 0$ and $R_O(T_S) = 1$ are both totally informative, corresponding to $T_S = O$ and $T_S = -O$, respectively, i.e., perfect learning (up to a possible sign error). In contrast, a risk value of $R_O(T_S) = 1/2$ implies that the hypothesis unitary matches the target no better than a typical random linear optical unitary.

The following theorem quantifies the expected risk for learning a linear optical unitary using the training set S , Eq. (37), in the limit of perfect learning.

Theorem 1: *Let O be distributed according to the normalized Haar measure on $\text{Orth}(2m)$, and let the training data S of cardinality $|S|$ be chosen uniformly from a compact connected set of m -mode coherent states, as defined*

in Eq. (37). Then

$$E_S\{E_O[R_O(T_S)]\} = \frac{1}{2} - \frac{|S|}{4m}. \quad (39)$$

Proof. For fixed S , simplify Eq. (38) to

$$\begin{aligned} R_O(T_S) &= \frac{1}{2} - \frac{1}{4m\sigma} \int d^{2m}x x T_S O^T x^T \pi(x) \\ &= \frac{1}{2} - \frac{\text{Tr} T_S O^T}{4m}. \end{aligned} \quad (40)$$

Under the assumption that the learning algorithm outputs T_S that agrees with O on the $|S|$ -dimensional subspace of \mathbb{R}^{2m} spanned by the training data (i.e., $w_j T = w_j O$ for the training data mean vectors w_j), we can write

$$T_S O^T = \begin{pmatrix} \mathbb{I} & 0 \\ 0 & Y \end{pmatrix}, \quad (41)$$

where $Y \in \text{Orth}(2m - |S|)$. Taking the expectation over O gives

$$\begin{aligned} E_O[R_O(T_S)] &= \int dO \left[\frac{1}{2} - \frac{\text{Tr} T_S O^T}{4m} \right] \\ &= \frac{1}{2} - \frac{|S|}{4m} - \int dY \frac{\text{Tr} Y}{4m} \\ &= \frac{1}{2} - \frac{|S|}{4m}. \end{aligned} \quad (42)$$

Because S is taken from a subset of \mathbb{R}^{2m} with no isolated points, one always obtains a set of $|S|$ linearly independent coherent states when S is sampled. Therefore, taking the expectation over S does not change the right-hand side of Eq. (42). ■

Theorem 1 shows that the generalization error for learning generic linear optic unitaries reduces linearly with the number of pairs of coherent states trained on, vanishing completely for $|S| = 2m$. (We stress that $|S|$ is the number of unique training pairs required to learn the unitary, not the total number, which, due to shot noise and the iterative optimization procedure, will be substantially larger.) This implies that the averaged coherent state cost in Eq. (17) can be approximated using only $2m$ training states when learning m -mode linear optical unitaries. More broadly, Theorem 1 can be viewed as a “classical” NFL theorem for CV systems.

It is possible to formulate a quasiclassical CV NFL theorem in which the training data consist of squeezed, rather than coherent, states. In this case we use a risk function that compares the actions of O and T_S on the phase space fluctuations of a compact set of centered, pure CV Gaussian states. We find that squeezing in general

inhibits the learning process. However, intriguingly, for this definition of the risk, the risk may be reduced by the training set size as a function of $|S|^2$ instead of $|S|$. This CV NFL theorem is discussed and proved in Appendix D.

We now show, similarly to the entanglement-assisted discrete-variable NFL theorem [34], that utilizing entangled training states can lower the expected risk. This improvement is achieved by modifying the training data set in Theorem 1, while keeping the risk (38) the same. Specifically, we now consider a training set

$$S = \{(|\psi_j^{\mathfrak{r}}\rangle, U \otimes \mathbb{I}_{\mathcal{R}} |\psi_j^{\mathfrak{r}}\rangle)\}_{j=1}^{|S|} \subset (\mathcal{H}_{\mathcal{X}} \otimes \mathcal{H}_{\mathcal{R}})^{\times 2|S|} \quad (43)$$

composed of $|S|$ pairs of m -mode *entangled coherent-Fock states* of the form

$$|\psi_j^{\mathfrak{r}}\rangle := \frac{1}{\sqrt{\mathfrak{r}}} \sum_{k=1}^{\mathfrak{r}} |w_k^{(j)}\rangle_{\mathcal{X}} \otimes |k\rangle_{\mathcal{R}}. \quad (44)$$

Here $\{|w_k^{(j)}\rangle_{\mathcal{X}}\}_{k=1}^{\mathfrak{r}}$ is a set of linearly independent coherent states acting on a system \mathcal{X} and $|k\rangle_{\mathcal{R}}$ denotes the k -th Fock state of an ancilla register \mathcal{R} . The positive integer \mathfrak{r} acts as an analogue of Schmidt rank in this context, although we note that the linearly independent mean vectors $w_k^{(j)}$ need not be approximately orthogonal, so \mathfrak{r} is not strictly related to the entanglement entropy. To use a precise term, \mathfrak{r} is equal to the exponential of the entropy of coherence [62] with respect to the orthonormal set $\{|w_k^{(j)}\rangle \otimes |k\rangle\}_{k=1}^{\mathfrak{r}}$ for any j . If $\|w_k^{(j)}\|$ is sufficiently large, $|\psi_j^{\mathfrak{r}}\rangle$ has entanglement entropy approximately equal to $\log_2 \mathfrak{r}$ with respect to the partition consisting of m CV modes \mathcal{X} and the CV register \mathcal{R} of the training set.

Analogously to the NFL for coherent state training above, we derive the following theorem on the expected risk.

Theorem 2: *Let O be distributed according to the normalized Haar measure on $\text{Orth}(2m)$, and let the training data S of cardinality $|S|$ consist of pairs of entangled coherent-Fock states as defined in Eqs. (43) and (44). Then*

$$E_S\{E_O[R_O(T_S)]\} = \frac{1}{2} - \frac{|S|\mathfrak{r}}{4m}. \quad (45)$$

Proof. As in the setting of Theorem 1, the objective is to learn the orthogonal matrix O corresponding to an m -mode linear optical unitary U . The assumption of perfect agreement of T_S and O on the training data set now corresponds to the condition $V_S \otimes \mathbb{I}_{\mathcal{R}} |\psi_j^{\mathfrak{r}}\rangle = U \otimes \mathbb{I}_{\mathcal{R}} |\psi_j^{\mathfrak{r}}\rangle$ for all j . Proceeding up to Eq. (40) in the same way as in the proof of Theorem 1, we now note that the assumption of perfect agreement on training data implies that $w_k^{(j)} O = w_k^{(j)} T_S$ for all k, j . Taking into account linear independence of the

mean vectors $w_k^{(j)}$ in \mathbb{R}^{2m} , this means that T_S and O are identical on a $\tau|S|$ -dimensional subspace of the phase space \mathbb{R}^{2m} . So $T_S O^T = \mathbb{I}_{\tau|S|} \oplus Y$ with $Y \in \text{Orth}(2m - \tau|S|)$ and, instead of Eq. (42) above, one gets

$$\begin{aligned} E_O[R_O(T_S)] &= \int dO \left[\frac{1}{2} - \frac{\text{Tr} T_S O^T}{4m} \right] \\ &= \frac{1}{2} - \frac{\tau|S|}{4m} - \int dY \frac{\text{Tr} Y}{4m} \\ &= \frac{1}{2} \left(1 - \frac{\tau|S|}{2m} \right). \end{aligned} \quad (46)$$

Again the expectation over training sets S of fixed cardinality $|S|$ is trivial when the mean vectors $w_k^{(j)}$ are chosen uniformly from some compact connected subset of \mathbb{R}^{2m} . ■

Theorem 2 shows that, for a fixed training data set size, increasing the parameter τ in the training data (for large $\|w_k^{(j)}\|$, this approximately corresponds to increasing the entanglement entropy of the training data) can reduce the generalization error. In this sense, entanglement could be seen to provide a “free-lunch.” However, as with all apparently free lunches, there are caveats. Namely, there may be a hidden cost in obtaining the entangled training data in the first place since entanglement is generally experimentally challenging to create and preserve. Thus, how “free” this lunch is will depend on the relative scarcity of training states and entanglement.

It is also important to note that the enhancement provided by entanglement here is less necessary than the enhancement found in the discrete-variable case. In the discrete-variable case an exponential number of training

pairs are required in the absence of entangled training data, whereas to learn linear optical unitaries, the number of unentangled pairs scales linearly in the number of modes.

Theorem 2 could be viewed as motivating a cost function of the form

$$C_{\text{ECFS}}^{(\tau,k)}(V, U) = 1 - \frac{1}{k} \sum_{j=1}^k |\langle \psi_j^\tau | V^\dagger U \otimes \mathbb{I}_{\mathcal{R}} | \psi_j^\tau \rangle|^2, \quad (47)$$

where the $|\psi_j^\tau\rangle$ are the entangled coherent-Fock states defined in Eq. (44). We note that this is a generalization of C_{ACS} in the sense that it reduces to C_{ACS} in the limit that $\tau = 1$. To learn a linear optical unitary, Theorem 2 implies that it suffices to use $k = 2m/\tau$ training pairs. One could also potentially use this cost to learn more general unitaries; however, Theorem 2 does not apply in that case and therefore one may need to use a significantly larger τk to minimize the generalization error.

In Appendix C we prove that Theorems 1 and 2 generalize to learning arbitrary Gaussian operations. We thus expect it to be possible to learn a single-mode Gaussian operation using a single entangled training pair ($\tau = 2$, $|S| = 1$), or two unentangled training pairs ($\tau = 1$, $|S| = 2$), but not a single unentangled training pair ($\tau = 1$, $|S| = 1$) since for the former the expected risk vanishes whereas the latter corresponds to a finite risk.

This is indeed supported by our numerical results shown in Fig. 6 where we optimize $C_{\text{ACS}}^{(1)}$ (corresponding to training on a single unentangled training state pair), $C_{\text{ACS}}^{(2)}$ (corresponding to training on two unentangled training state pairs), and $C_{\text{ECFS}}^{(2,1)}$ (corresponding to training on a single entangled training state pair) using the same variational framework set out in Sec. IV. We find that, while




FIG. 6. Learning Gaussian operations using coherent states and entangled coherent-Fock states. Cost as a function of iteration for learning a Gaussian operation. In the left and middle plots we use C_{ACS} with a single training pair ($k = 1$) and two training pairs ($k = 2$), respectively. In the right plot we optimize $C_{\text{ECFS}}^{(\tau=2,k=1)}(V, U)$ that is defined in Eq. (47) from one training state of the form (44). In all cases the coherent states trained on had a random energy up to a maximum of 1 and a random phase in the range 0 to 2π . To quantify the quality of the optimization, we take the optimal parameters obtained at each iteration of the optimization algorithm and plot both the Hilbert-Schmidt cost C_{HST} (red) and the errors in the individual optimized parameters in arbitrary units (dotted). The dashed and dash-dot red lines show the Hilbert-Schmidt cost C_{HST} on the first 50 Fock states and on the first 5 Fock states, respectively.

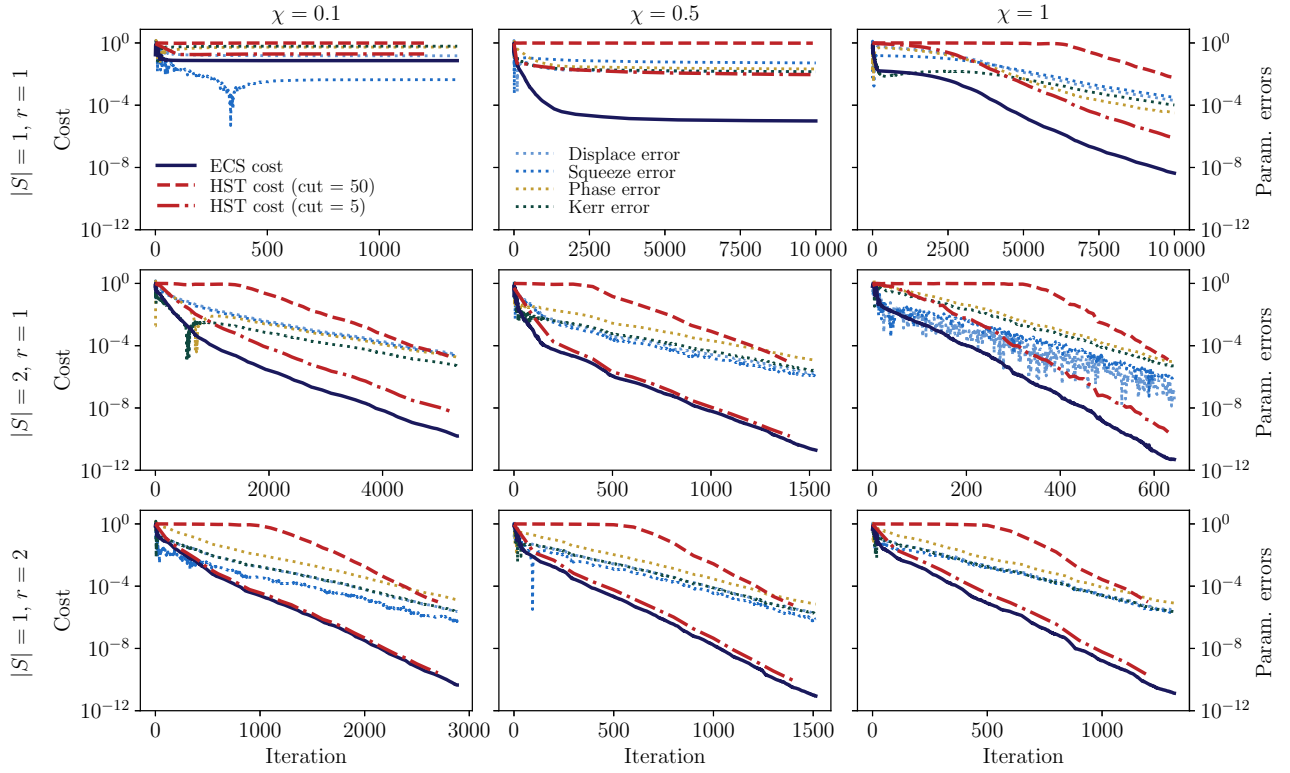


FIG. 7. Learning Kerr nonlinearities using coherent states and entangled coherent-Fock states. Cost as a function of iteration for learning a $\chi = 0.1$ (top) and $\chi = 0.5$ (bottom) nonlinearity. In the left and middle columns we use C_{ACS} with a single training pair ($k = 1$) and two training pairs ($k = 2$), respectively. In the right-hand column we use C_{ECFS} with a single entangled training pair ($k = 1$ and $\tau = 2$). In all cases the basic coherent state training states have a random energy up to a maximum of 1 and a random phase in the range 0 to 2π . To quantify the quality of the optimization, we take the optimal parameters obtained at each iteration of the optimization algorithm and plot both the Hilbert-Schmidt cost C_{HST} (red) and the errors in the individual optimized parameters in arbitrary units (dotted). The dashed red line and dash-dot red lines show the truncated Hilbert-Schmidt cost C_{HST} on the first 50 Fock states and on the first 5 Fock states, respectively.

it is possible to minimize $C_{ACS}^{(1)}$, this does not correspond to the Gaussian operation being successfully learnt. This is shown by the large learning errors, as measured by the truncated Hilbert-Schmidt test cost, which quantifies the average error over all possible input states, and individual parameter errors, in the left-hand panel of Fig. 6. Conversely, as shown in the middle and right-hand panels of Fig. 6, when using entangled training data or multiple training states, the learning errors are iteratively minimized as the cost is minimized.

In Fig. 7 we present analogous results for the learning of a weak single-mode Kerr nonlinearity. Specifically, as shown in Fig. 7 we find that a single-mode ($m = 1$) Kerr nonlinearity of $\chi = 0.1$ and $\chi = 0.5$ can be learnt using either a single entangled training pair ($\tau = 2$, $|S| = 1$), or two unentangled training pairs ($\tau = 1$, $|S| = 2$), but not a single unentangled training pair ($\tau = 1$, $|S| = 1$).

B. Learning arbitrary unitaries and motivation of compiling cost functions

Theorems 1 and 2 concern learning linear optical unitaries. The question remains whether similar

entanglement-assisted NFL theorems can be derived for learning arbitrary CV unitaries.

To answer this question, it is useful to recall the discrete-variable (i.e., finite-dimensional) entanglement-assisted quantum NFL theorem in Ref. [34]. Specifically, the analog of Theorem 2 takes the form

$$E_S\{E_U[R_U(V_S)]\} = 1 - \frac{\tau^2|S|^2 + d + 1}{d(d + 1)}, \quad (48)$$

where U is the target unitary, V_S is the output of the learning algorithm on entangled training states in S , τ is the Schmidt rank of the training data states, and the risk is

$$R_U(V_S) := \frac{1}{4} \int d\psi \|U|\psi\rangle\langle\psi|U^\dagger - V_S|\psi\rangle\langle\psi|V_S^\dagger\|_1^2. \quad (49)$$

In Eq. (49), the integral is over all pure states according to the Haar measure.

A continuous-variable NFL theorem cannot be derived that is strictly analogous to Eq. (48) in the discrete-variable setting because there is no Haar measure over the unitary

group in $B(\mathcal{H})$ for infinite dimensional \mathcal{H} . On the other hand, one is often only interested in the action of the target unitary U on Fock states only up to a finite cutoff. For example, recent proposals for efficient updates and derivatives of Gaussian gates in parameterized CV circuits utilize cutoff recursion relations for the Fock matrix elements of the gates [5].

Equation (48) implies that a single full rank state, i.e., a state with rank d , can be used to fully learn a unitary of rank d . Thus, the truncated TMSSs defined in Eq. (8) can be used to learn arbitrary d -dimensional unitaries without incurring a generalization error. Taking the limit that d tends to infinity, this implies that the Loschmidt echo TMSS cost can be used to learn arbitrary CV unitaries, thus further motivating its use.

VI. DISCUSSION

In this work we establish a framework for quantum compiling in continuous-variable systems. We start by motivating the TMSS cost (both the Loschmidt echo and ricocheted variants) and the averaged coherent state cost as natural CV analogues of the Hilbert-Schmidt state cost. Our numerical implementations demonstrate the successful learning of single-mode Gaussian operations, a generalized beam splitter operation, and Kerr nonlinearities using these costs.

We subsequently show how these costs may be alternatively motivated via a series of increasingly general “(no-)free lunch” theorems. Firstly, the NFL theorem for Gaussian operations using coherent state mean vector training data establishes that it is possible to perfectly learn an m -mode Gaussian operation by training on only $2m$ coherent states. This implies that it is possible to learn arbitrary Gaussians by training on an approximation of the averaged coherent state cost using only $2m$ coherent states. Next, the NFL theorem for Gaussian operations using entangled coherent-Fock states both shows how entanglement may be used to reduce the amount of training data required to learn Gaussian operations and motivates an alternative entanglement-enhanced cost function for compiling that makes use of entangled coherent-Fock states. Finally, we argue that taking the continuum limit of the discrete-variable entanglement-enhanced NFL theorem implies that to learn an arbitrary unitary on a single training state requires a full rank state. This motivates training using the TMSS cost.

It is worth highlighting that these (no-)free-lunch theorems quantify the number of *different* training pairs required to learn a unitary in the ideal case of perfect training. That is, they *do not* give the *total* number of copies of training pairs that are required to learn the unitary. Indeed, given shot noise, a large number of copies of each pair will in fact be required to evaluate the cost. More generally, training may be imperfect not only due to shot noise

but also hardware noise or the presence of barren plateaus or local minima in the training cost function landscape. A valuable extension would be to generalize the theorems to account for imperfect learning.

It would also be interesting to derive further NFL theorems for alternative classes in training data. For example, one might be concerned with learning a unitary U from homodyne or heterodyne detection data, in which case a risk function could be defined in terms of the difference in the expected quadrature vector of the output state for the hypothesis and target unitaries. General unitary learning protocols based on other CV measurement-motivated risk functions, such as those associated with CV distinguishability norms [63,64], are expected to have associated NFL theorems and quantum compiling protocols that are adapted to the measurement class under consideration.

We further note that two-mode-squeezed states are not the only choice of state to saturate the entanglement-enhanced NFL bound for arbitrary unitaries. One could alternatively use any full rank state, such as cluster states. A finite energy CV cluster state is defined by $|\text{CL}_r\rangle = e^{iq\otimes q}[S(-r)|0\rangle]^{\otimes 2}$, where q is the single-mode position quadrature and $S(r)$ is the unitary squeezing operator. The state $|\text{CL}_r\rangle$ limits to the well-known CV cluster state for $r \rightarrow \infty$ [65]. One could use $|\text{CL}_r\rangle$ to define a faithful cost function analogous to the Loschmidt echo and ricocheted TMSS costs.

As quantum hardware develops, the CV quantum compiling algorithms we have presented here are expected to find use optimizing short-depth CV quantum circuits, thereby aiding the implementation of larger scale quantum algorithms. Furthermore, we envision that tuning CV quantum resources such as intensity or squeezing could allow one to implement our CV quantum compiling algorithms in a noise resistant way. For example, results of Ref. [52] indicate that sublinear scaling (with mode number m) of the coherent state intensity and number of quantum-limited attenuator layers does not induce barren plateaus in cost functions such as C_{ACSE} when restricted to linear optical unitaries. More generally, we are excited by the idea that these quantum compilation algorithms may be used to study the optical properties of new materials. It would be interesting to explore whether these algorithms could be combined with metalearning strategies to actively design new materials with desirable properties such as controllable squeezing amplitudes or nonlinearities.

ACKNOWLEDGMENTS

The authors thank Kunal Sharma and Patrick Coles for helpful discussions. T.V. acknowledges support from the LDRD program at LANL. Z.H. acknowledges support and A.S. acknowledges initial support from the LANL ASC Beyond Moore’s Law project. This material is based upon work supported by the U.S. Department of Energy,

Office of Science, National Quantum Information Science Research Centers, Quantum Science Center (AS).

APPENDIX A: COMPARISON OF LOSCHMIDT ECHO (10) AND RICOCHETED (13) TMSS COSTS

We show how the difference between cost functions (10) and (13) depends on the squeezing parameter r and the rank of the variational ansatz V . Consider the truncated two-mode-squeezed state in Eq. (8) and variational ansatz V such that the rank of V is τ . For V distributed with respect to the Haar measure on the unitary group $U(\tau)$, one can see that the states $V^\dagger \otimes \mathbb{I} |\psi_{\text{TMSS}}^\tau(r)\rangle$ and $\mathbb{I} \otimes V^* |\psi_{\text{TMSS}}^\tau(r)\rangle$ are nearly equal in expectation for large r . Specifically, the expected modulus of the inner product of these states is given by

$$\begin{aligned} & E_V \{ |(V^\dagger \otimes \mathbb{I} |\psi_{\text{TMSS}}^\tau(r)\rangle, \mathbb{I} \otimes V^* |\psi_{\text{TMSS}}^\tau(r)\rangle)| \} \\ &= \frac{1 - \tanh^2 r}{1 - \tanh^{2\tau} r} E_V \left(\sum_{\ell, \ell'=0}^{\tau} \tanh^{\ell+\ell'} r |V_{\ell, \ell'}|^2 \right) \\ &= \frac{1}{\tau} \frac{1 - \tanh^2 r}{1 - \tanh^{2\tau} r} \left(\frac{1 - \tanh^\tau r}{1 - \tanh r} \right)^2 \\ &= \frac{1}{\tau} \frac{1 + \tanh r}{1 + \tanh^\tau r} \frac{1 - \tanh^\tau r}{1 - \tanh r} \\ &\sim \tanh^{\tau-1} r \quad \text{for } r \rightarrow \infty. \end{aligned} \quad (\text{A1})$$

For fixed τ , the $r \rightarrow \infty$ limit is 1. Furthermore, even if τ is increased by a multiplicative factor $\tau \mapsto \lambda\tau$, i.e., the unitaries considered are in $U(\lambda\tau)$, the value of the expectation remains close to 1 if one simply adjusts the squeezing r according to $r \mapsto r + \ln \lambda$. To see this, just expand the asymptotic function in Eq. (A1) with respect to the small number e^{-2r} to get $\tanh^{\tau-1} r \sim 1 - 2(\tau-1)e^{-2r}$ for large r .

APPENDIX B: FAITHFULNESS OF $\tilde{C}_{\text{R-TMSS}_r}$

To prove the faithfulness of $\tilde{C}_{\text{R-TMSS}_r}$, we start by showing that it can be written in terms of the inner product $(V, U)_{\rho_\beta^{\otimes m}} := \text{Tr} \sqrt{\rho_\beta^{\otimes m}} U \sqrt{\rho_\beta^{\otimes m}} V^\dagger$, where $\rho_\beta \propto \sum_{n=0}^{\infty} e^{-\beta n} |n\rangle\langle n|$ is a single-mode thermal state with inverse temperature β . This inner product appears in the theory of generalized conditional expectations [66] and quantum relative entropies [67]. Despite its complicated appearance, $(V, U)_{\rho_\beta^{\otimes m}}$ is actually efficiently computable using pure entangled Gaussian state preparation. To demonstrate this fact, consider m copies of a two-mode-squeezed state prepared in mode pairs (A_j, B_j) , $j = 1, \dots, m$:

$$|\psi_{\text{TMSS}}^m(r)\rangle \propto \bigotimes_{j=1}^m \sum_{\ell=0}^{\infty} (\tanh r)^\ell |\ell\rangle_{A_j} \otimes |\ell\rangle_{B_j} \quad (\text{B1})$$

with squeezing parameter r satisfying $-2 \ln \tanh r = \beta$. It follows that

$$\begin{aligned} (V, U)_{\rho_\beta^{\otimes m}} &= (1 - e^{-\beta})^m \sum_{\ell, \ell'} e^{-\beta \|\ell + \ell'\|_1 / 2} U_{\ell, \ell'} (V^\dagger)_{\ell', \ell} \\ &= (1 - \tanh^2 r)^m \sum_{\ell, \ell'} (\tanh r)^{\|\ell + \ell'\|_1} U_{\ell, \ell'} V_{\ell', \ell}^* \\ &= \text{Tr} |\psi_{\text{TMSS}}^m(r)\rangle \langle \psi_{\text{TMSS}}^m(r)|^{\otimes m} U_A \otimes V_B^*, \end{aligned} \quad (\text{B2})$$

where the sums over ℓ, ℓ' are over $\mathbb{Z}_{\geq 0}^{\times m}$. It follows from definition (13) that $C_{\text{R-TMSS}_r}(V, U) = 1 - |(V, U)_{\rho_\beta^{\otimes m}}|^2$ and from Eq. (14) that $\tilde{C}_{\text{R-TMSS}_r}(V, U) = 1 - |(V, U)_{\rho_\beta^{\otimes m}}|^2 / |(V, V)_{\rho_\beta^{\otimes m}}|^2 |(U, U)_{\rho_\beta^{\otimes m}}|^2$. The facts that $(X, Y)_{\rho_\beta^{\otimes m}}$ is linear in Y , conjugate linear in X , $(X, X)_{\rho_\beta^{\otimes m}} \in \mathbb{R}$ (with value 0 if and only if $X = 0$) are clear. Therefore, the Cauchy-Schwarz inequality

$$|(V, U)_{\rho_\beta^{\otimes m}}|^2 \leq |(V, V)_{\rho_\beta^{\otimes m}}| |(U, U)_{\rho_\beta^{\otimes m}}| \quad (\text{B3})$$

holds. It implies the faithfulness of Eq. (14), i.e., $\tilde{C}_{\text{R-TMSS}_r}(V, U) = 0$ if and only if $U = e^{i\phi} V$ for some $\phi \in [0, 2\pi)$.

APPENDIX C: CV NFL THEOREMS FOR GAUSSIAN OPERATIONS

Here, we show that Theorems 1 and 2 can be generalized to learning arbitrary Gaussian operations.

In Corollary 1 below, the target unitary U is associated with $L \in \text{Sp}(2m, \mathbb{R})$ by $U^\dagger R U = R L$ and the learning algorithm outputs $T_S \in \text{Sp}(2m, \mathbb{R})$ when given training set S in Eq. (37). A $2m \times 2m$ real matrix L is symplectic if and only if $L^T \Delta L = \Delta$, where

$$\Delta = \begin{pmatrix} 0 & 1 \\ -1 & 0 \end{pmatrix}^{\oplus m}$$

is the standard symplectic form on \mathbb{R}^{2m} .

Corollary 1: For any $m \in \mathbb{N}$, define $\mathcal{G}^{(2m)} := \{O_1 Z O_2\}$, where $O_1, O_2 \in \text{Orth}(2m) \cap \text{Sp}(2m, \mathbb{R})$ and $Z = \bigoplus_{j=1}^m \text{diag}(z_j, z_j^{-1})$ with $z_j \in \mathbb{R}_+$. Let S be the training data (37) with $|S| \equiv 0 \pmod{2}$, and let $L \in \text{Sp}(2m, \mathbb{R})$. If $R_L(T_S)$ is the risk (38) and $\mathcal{G}^{(2m)}$ is equipped with the probability measure $dO_1 dO_2 \mu(dz)$, with dO_1 and dO_2 Haar measures and $\mu(dz)$ a probability measure on \mathbb{R}^m , then $E_S\{E_L[R_L(T_S)]\}$ is given by Eq. (39).

Proof. Every symplectic matrix L is in $\mathcal{G}^{(2m)}$ due to the Bloch-Messiah decomposition [68]. From Eq. (40), it follows that $R_L(T_S) = \frac{1}{2} - \text{Tr} T_S L^T / 4m$. Since T_S and L are

assumed to agree on the subspace of \mathbb{R}^{2m} spanned by the training data, one can write

$$T_S L^T = \begin{pmatrix} \mathbb{I}_{|S|} & A \\ B & Y \end{pmatrix}.$$

But $T_S L^T \in \text{Sp}(2m, \mathbb{R})$ implies that

$$\Delta_{|S|} + B^T \Delta_{2m-|S|} B = \Delta_{|S|}, \quad (\text{C1a})$$

$$Y^T \Delta_{2m-|S|} Y + A^T \Delta_{|S|} A = \Delta_{2m-|S|}, \quad (\text{C1b})$$

where $\Delta = \Delta_{|S|} \oplus \Delta_{2m-|S|}$. Equation (C1a) implies that $B = 0$ from nondegeneracy of the symplectic form. Let V be the Gaussian unitary that acts on the canonical operators as $V^\dagger R V = R T_S L^T$. Then the action of V on the $2m$ vector of operators $(0, \dots, 0, \tilde{R})$ with $\tilde{R} := (q_{|S|+1}, p_{|S|+1}, \dots, q_m, p_m)$ is $(\tilde{R} B, \tilde{R} Y) = (0, \dots, 0, \tilde{R} Y)$, so the unitary invariance of the canonical commutation relation implies that $Y \in \text{Sp}(2m - |S|)$. It then follows from Eq. (C1b) that $A = 0$. One concludes that

$$E_L[R_L(T_S)] = \frac{1}{2} - \frac{|S|}{4m} - E_Y \left(\frac{\text{Tr} Y}{4m} \right), \quad (\text{C2})$$

where E_Y is taken with respect to the measure on Y induced by the measure on $\mathcal{G}^{(2m)}$. Since Y is a symplectic matrix, it can be written $W_1 F W_2$ with $W_1, W_2 \in \text{Orth}(2m - |S|)$. The fact that W_1 and W_2 are independent and distributed according to the Haar measure follows from restricting the Haar measure on O_1 and O_2 . Therefore, the expectation over Y in Eq. (C2) is zero. ■

An entirely equivalent argument can be used to generalize Theorem 2 to learning Gaussian operations.

APPENDIX D: CV NFL THEOREM WITH GAUSSIAN TRAINING DATA

A centered Gaussian state is a Gaussian state that satisfies $\langle R \rangle = 0$ and, therefore, is uniquely defined by its covariance matrix $\Sigma_{i,j} = \frac{1}{2} \langle [R_i, R_j]_+ \rangle$, which is a positive $2m \times 2m$ matrix. For example, the only centered coherent state is the vacuum $\Sigma_{|0\rangle} = \text{diag}(\frac{1}{2}, \dots, \frac{1}{2})$. For examples with entanglement, the two-mode-squeezed states and CV cluster states are pure, centered Gaussian states with $m = 2$. Let $|\psi\rangle$ be a centered Gaussian state, and let U be a linear optical unitary that satisfies $U^\dagger R U = R O$. Then $\Sigma_{U|\psi\rangle} = O^T \Sigma_{|\psi\rangle} O$. Instead of training with coherent state mean vectors, consider now linearly independent training data $S = \{(\Sigma^{(j)}, O^T \Sigma^{(j)} O)\}_{j=1}^{|S|}$, where $\Sigma^{(j)}$ is the covariance matrix of an m -mode pure, centered Gaussian state satisfying $\text{rank} \frac{\mathbb{I}_{2m}}{2} - \Sigma^{(j)} = 2$ for each j , i.e., the state is squeezed only in one phase space direction. We consider

the risk function

$$R_O(T_S) = \frac{1}{m(\log D)^m} \int_{\Omega} d\Sigma \int \|\mathbb{I}_{2m}^T \Sigma T_S - O^T \Sigma O\|_2^2, \quad (\text{D1})$$

where the integral $d\Sigma$ is taken over a compact subset Ω of covariance matrices that satisfy $\|\Sigma\| \leq \frac{D}{2}$ and $|\det 2\Sigma| = 1$. Physically, Ω is the set of covariance matrices of pure Gaussian states with maximal squeezing parameter $r = \frac{1}{2} \log D$. We assume that the learning algorithm outputs an orthogonal matrix T_S such that $T_S^T \Sigma^{(j)} T_S = O^T \Sigma^{(j)} O$ for all j , i.e., the algorithm produces perfect agreement with the target on the training data set. With the cost function (D1), Theorem 3 shows that the expected risk is reduced by a function scaling as $|S|^2$ instead of $|S|$ in Theorem 1.

Theorem 3: *Let O be distributed according to the normalized Haar measure on $\text{Orth}(2m)$, and let the training data S of cardinality $|S|$ be chosen uniformly from a compact subset of m -mode pure, centered Gaussian states satisfying the rank condition above. Then, for the risk function (D1),*

$$E_S\{E_O[R_O(T_S)]\} = \frac{D^2 - D^{-2}}{4 \log D} \left(1 - \frac{1}{2(m+1)} \right) - \frac{(D - D^{-1})^2 (|S|^2 + 1)}{8m(\log D)^2} + \mathcal{O}(m^{-2}). \quad (\text{D2})$$

Proof. The integral defining the risk (D1) is over a compact set Ω of pure, centered Gaussian states that have squeezing parameters with magnitude uniformly distributed between $r = -\frac{1}{2} \log D$ and $r = \frac{1}{2} \log D$, where $D > 1$. Specifically, the covariance matrices Σ appearing in the integral have the form

$$\Sigma \in \left\{ W^T \text{diag} \left(\frac{e^{-2r_1}}{2}, \frac{e^{2r_1}}{2}, \dots, \frac{e^{-2r_m}}{2}, \frac{e^{2r_m}}{2} \right) W : W \in \text{Orth}(2m), r_j \in \left[-\frac{1}{2} \log D, \frac{1}{2} \log D \right] \right\}. \quad (\text{D3})$$

For calculating the expected risk function, it is advantageous to use the **vec** functor. For a matrix $A \in \text{End}(\mathbb{R}^{2m})$,

$$\text{vec} A = \sum_{i=1}^{2m} \sum_{j=1}^{2m} A_{i,j} e_i \otimes e_j \in \mathbb{R}^{4m^2}, \quad (\text{D4})$$

where $\{e_j\}_{j=1}^{2m}$ is an orthonormal basis of \mathbb{R}^{2m} . For example, $\text{vec}(W^T A W) = (W \otimes W) \text{vec}(A)$. Also, **vec** is an isometry from $\text{End}(\mathbb{R}^{2m})$ as a finite-dimensional Hilbert space with Hilbert-Schmidt inner product to \mathbb{R}^{4m^2} as a Hilbert space with Euclidean inner product: $\|A\|_2^2 = \|\text{vec} A\|^2$. The risk (D1) becomes

$$\begin{aligned}
R_O(T_S) &= \frac{1}{m(\log D)^m} \int_{-(1/2)\log D}^{(1/2)\log D} dr_1 \cdots \int_{-(1/2)\log D}^{(1/2)\log D} dr_m \int dW \\
&\times \left[\|T_S^T W^T \text{diag}\left(\frac{e^{-2r_1}}{2}, \frac{e^{2r_1}}{2}, \dots, \frac{e^{-2r_m}}{2}, \frac{e^{2r_m}}{2}\right) W T_S - O^T W^T \text{diag}\left(\frac{e^{-2r_1}}{2}, \frac{e^{2r_1}}{2}, \dots, \frac{e^{-2r_m}}{2}, \frac{e^{2r_m}}{2}\right) W O\|_2^2 \right] \\
&= \frac{1}{4m(\log D)^m} \int d\mathbf{r} dW \| (T_S \otimes T_S - O \otimes O)(W \otimes W) \sum_{j=1}^m (e^{-2r_j} e_{2j-1} \otimes e_{2j-1} + e^{2r_j} e_{2j} \otimes e_{2j}) \|^2 \quad (\text{D5})
\end{aligned}$$

with dW the normalized Haar measure over $\text{Orth}(2m)$, where the second equality follows from using the isometric property of vec . In the last equality, we also shortened the integral notation. Expanding the square of the Euclidean distance gives a sum of two integrals:

$$\begin{aligned}
R_O(T_S) &= \frac{2}{4m(\log D)^m} \int d\mathbf{r} \sum_{j=1}^m (e^{-4r_j} + e^{4r_j}) - \frac{2}{4m(\log D)^m} \int d\mathbf{r} \\
&\times \int dW \sum_{k,j=1}^m [e^{-2(r_j+r_k)} (e_{2j-1}^T W^T T_S^T O W e_{2k-1})^2 + e^{-2(r_j-r_k)} (e_{2j-1}^T W^T T_S^T O W e_{2k})^2 \\
&+ e^{2(r_j-r_k)} (e_{2j}^T W^T T_S^T O W e_{2k-1})^2 + e^{2(r_j+r_k)} (e_{2j}^T W^T T_S^T O W e_{2k})^2]. \quad (\text{D6})
\end{aligned}$$

The integral in the first line evaluates to $(D^2 - D^{-2})/4 \log D$. In the second integral, it is useful to break up the sum over j, k to the $j = k$ and $j \neq k$ parts. Then we use the following lemma.

Lemma 1: *Let V be a real vector space with orthonormal basis $\{e_i\}_{i=1}^{2m}$, and let $W, L \in \text{Orth}(2m)$. Then, for any $i, j \in \{1, \dots, 2m\}$ with $j \neq i$,*

$$\begin{aligned}
\int dW (e_i^T W^T L W e_i)^2 &= \frac{2m + \text{Tr} L^2 + (\text{Tr} L)^2}{4m^2 + 4m}, \\
\int dW (e_i^T W^T L W e_j)^2 &= \sum_{r=1}^{2m} \frac{L_{r,r}^2}{4m^2 + 4m} \\
&+ \sum_{r \neq r'} \frac{L_{r,r} L_{r',r'} (2m+1)}{4m(m+1)(2m-1)} \\
&\rightarrow \frac{(\text{Tr} L)^2}{4m^2 + 4m} \quad \text{as } m \rightarrow \infty. \quad (\text{D7})
\end{aligned}$$

The proof of the lemma involves integration over the orthogonal group with respect to the Haar measure [69]. We apply the lemma with $L = T_S^T O$, using the first integral from the lemma exactly to evaluate the $j = k$ part of the second integral in Eq. (D6) and using the second integral from the lemma in its asymptotic form to evaluate the $j \neq k$ part of the second integral in Eq. (D6). The result for

integration over W is

$$\begin{aligned}
R_O(T_S) &= \frac{D^2 - D^{-2}}{4 \log D} \\
&- \frac{(D^2 - D^{-2})[2m + \text{Tr} T_S^T O T_S^T O + (\text{Tr} T_S^T O)^2]}{16m(m+1) \log D} \\
&- \frac{(D - D^{-1})^2 (\text{Tr} T_S^T O)^2}{8m(\log D)^2}. \quad (\text{D8})
\end{aligned}$$

The assumption that T_S and O agree on the training dataset S is now taken into account. Recall that the training covariance matrices $\Sigma^{(j)}$ are associated with distinct directions in \mathbb{R}^{2m} . Therefore, we can write $T_S^T O = \mathbb{I}_{|S|} \oplus Y$ with $Y \in \text{Orth}[(2m - |S|)]$. Note that $\int dY \text{Tr} Y = 0$, $\int dY \text{Tr} Y^2 = 1$, and $\int dY (\text{Tr} Y)^2 = 1$. From this, it follows that $E_O[(\text{Tr} T_S^T O)^2] = E_Y[(|S| + \text{tr} Y)^2] = |S|^2 + 1$ and $E_O\{\text{Tr}[(T_S^T O)^2]\} = E_Y(|S| + Y^2) = |S| + 1$. Applying these to Eq. (D8) gives

$$\begin{aligned}
E_O[R_O(T)] &= \frac{D^2 - D^{-2}}{4 \log D} - \frac{(D^2 - D^{-2})(2m + |S|^2 + |S| + 2)}{16m(m+1) \log D} \\
&- \frac{(D - D^{-1})^2 (|S|^2 + 1)}{8m(\log D)^2} + \mathcal{O}(m^{-2}), \quad (\text{D9})
\end{aligned}$$

where the $\mathcal{O}(m^{-2})$ comes from the asymptotic in Eq. (D7). Absorbing the remaining $\mathcal{O}(m^{-2})$ terms and carrying out the trivial average over the finite set S results in Eq. (D2). ■

- [1] G. R. Steinbrecher, J. P. Olson, D. Englund, and J. Carolan, Quantum optical neural networks, *npj Quantum Inf.* **5**, 60 (2019).
- [2] Yichen Shen, Nicholas C. Harris, Scott Skirlo, Mihika Prabhu, Tom Baehr-Jones, Michael Hochberg, Xin Sun, Shijie Zhao, Hugo Larochelle, Dirk Englund, and Marin Soljacic, Deep learning with coherent nanophotonic circuits, *Nat. Photonics* **11**, 441 (2017).
- [3] N. C. Harris, J. Carolan, D. Bunandar, M. Prabhu, M. Hochberg, T. Baehr-Jones, M. L. Fanto, A. M. Smith, C. C. Tison, P. M. Alsing, and D. Englund, Linear programmable nanophotonic processors, *Optica* **5**, 1623 (2018).
- [4] Hoi-Kwan Lau, Raphael Pooser, George Siopsis, and Christian Weedbrook, Quantum Machine Learning Over Infinite Dimensions, *Phys. Rev. Lett.* **118**, 080501 (2017).
- [5] Filippo M. Miatto and Nicolás Quesada, Fast optimization of parametrized quantum optical circuits, *Quantum* **4**, 366 (2020).
- [6] Leonardo Banchi, Nicolás Quesada, and Juan Miguel Arrazola, Training Gaussian boson sampling distributions, *Phys. Rev. A* **102**, 012417 (2020).
- [7] M. Cerezo, Andrew Arrasmith, Ryan Babbush, Simon C. Benjamin, Suguru Endo, Keisuke Fujii, Jarrod R. McClean, Kosuke Mitarai, Xiao Yuan, Lukasz Cincio, and Patrick J. Coles, Variational quantum algorithms, *Nat. Rev. Phys.* **3**, 625 (2021).
- [8] Kishor Bharti, Alba Cervera-Lierta, Thi Ha Kyaw, Tobias Haug, Sumner Alperin-Lea, Abhinav Anand, Matthias Degroote, Hermann Heimonen, Jakob S. Kottmann, Tim Menke, Wai-Keong Mok, Sukin Sim, Leong-Chuan Kwek, and Alán Aspuru-Guzik, Noisy intermediate-scale quantum (NISQ) algorithms, arXiv preprint [ArXiv:2101.08448](https://arxiv.org/abs/2101.08448) (2021).
- [9] Alberto Peruzzo, Jarrod McClean, Peter Shadbolt, Man-Hong Yung, Xiao-Qi Zhou, Peter J. Love, Alán Aspuru-Guzik, and Jeremy L. O'Brien, A variational eigenvalue solver on a photonic quantum processor, *Nat. Commun.* **5**, 4213 (2014).
- [10] J. Carolan, M. Mohseni, J. P. Olson, M. Prabhu, C. Chen, D. Bunandar, M. Y. Niu, N. C. Harris, F. N. C. Wong, M. Hochberg, S. Lloyd, and D. Englund, Variational quantum unsampling on a quantum photonic processor, *Nat. Phys.* **16**, 322 (2020).
- [11] Victor V. Albert, Kyungjoo Noh, Kasper Duivenvoorden, Dylan J. Young, R. T. Brierley, Philip Reinhold, Christophe Vuillot, Linshu Li, Chao Shen, S. M. Girvin, Barbara M. Terhal, and Liang Jiang, Performance and structure of single-mode bosonic codes, *Phys. Rev. A* **97**, 032346 (2018).
- [12] T. J. Volkoff, Nonclassical properties and quantum resources of hierarchical photonic superposition states, *J. Exp. Theor. Phys.* **121**, 770 (2015).
- [13] Daiqin Su, Robert Israel, Kunal Sharma, Haoyu Qi, and Ish Dhand, and Kamil Brádler, Error mitigation on a near-term quantum photonic device, *Quantum* **5**, 452 (2021).
- [14] Sumeet Khatri, Ryan LaRose, Alexander Poremba, Lukasz Cincio, Andrew T. Sornborger, and Patrick J. Coles, Quantum-assisted quantum compiling, *Quantum* **3**, 140 (2019).
- [15] Juan Miguel Arrazola, Thomas R. Bromley, Josh Izaac, Casey R. Myers, Kamil Brádler, and Nathan Killoran, Machine learning method for state preparation and gate synthesis on photonic quantum computers, *Quantum Sci. Technol.* **4**, 024004 (2019).
- [16] Kunal Sharma, Sumeet Khatri, M. Cerezo, and Patrick J. Coles, Noise resilience of variational quantum compiling, *New J. Phys.* **22**, 043006 (2020).
- [17] C. L. Degen, F. Reinhard, and P. Cappellaro, Quantum sensing, *Rev. Mod. Phys.* **89**, 035002 (2017).
- [18] Gus Gutoski and Nathaniel Johnston, Process tomography for unitary quantum channels, *J. Math. Phys.* **55**, 032201 (2014).
- [19] Charles H. Baldwin, Amir Kalev, and Ivan H. Deutsch, Quantum process tomography of unitary and near-unitary maps, *Phys. Rev. A* **90**, 012110 (2014).
- [20] T. J. Volkoff, Optimal and near-optimal probe states for quantum metrology of number-conserving two-mode bosonic hamiltonians, *Phys. Rev. A* **94**, 042327 (2016).
- [21] T. J. Volkoff and C. M. Herdman, Generating accessible entanglement in bosons via pair-correlated tunneling, *Phys. Rev. A* **100**, 022331 (2019).
- [22] C. Cirstoiu, Z. Holmes, J. Iosue, L. Cincio, P. J. Coles, and A. Sornborger, Variational fast forwarding for quantum simulation beyond the coherence time, *Jnpj Quantum Inf.* **6**, 82 (2020).
- [23] Joe Gibbs, Kaitlin Gili, Zoë Holmes, Benjamin Commeau, Andrew Arrasmith, Lukasz Cincio, Patrick J. Coles, and Andrew Sornborger, Long-time simulations with high fidelity on quantum hardware, arXiv preprint [ArXiv:2102.04313](https://arxiv.org/abs/2102.04313) (2021).
- [24] Michael R. Geller, Zoë Holmes, Patrick J. Coles, and Andrew Sornborger, Experimental quantum learning of a spectral decomposition, arXiv preprint [ArXiv:2104.03295](https://arxiv.org/abs/2104.03295) (2021).
- [25] Timjan Kalajdzievski, Christian Weedbrook, and Patrick Rebentrost, Continuous-variable gate decomposition for the Bose-Hubbard model, *Phys. Rev. A* **97**, 062311 (2018).
- [26] Timjan Kalajdzievski and Juan Miguel Arrazola, Exact gate decompositions for photonic quantum computing, *Phys. Rev. A* **99**, 022341 (2019).
- [27] R. Annabestani, B. Gupt, and B. R. Bardhan, Towards quantum simulation of spin systems using continuous variable quantum devices, arXiv preprint [ArXiv:2009.09455](https://arxiv.org/abs/2009.09455) (2020).
- [28] M. Horodecki, P. Horodecki, and R. Horodecki, General teleportation channel, singlet fraction, and quasidistillation, *Phys. Rev. A* **60**, 1888 (1999).
- [29] M. A. Nielsen, A simple formula for the average gate fidelity of a quantum dynamical operation, *Phys. Lett. A* **303**, 249 (2002).
- [30] Junaid Aasi, J. Abadie, B. P. Abbott, Richard Abbott, T. D. Abbott, M. R. Abernathy, Carl Adams, Thomas Adams, Paolo Addesso, and R. X. Adhikari *et al.*, Enhanced sensitivity of the LIGO gravitational wave detector by using squeezed states of light, *Nat. Photonics* **7**, 613 (2013).
- [31] B. J. Lawrie, P. D. Lett, A. M. Marino, and R. C. Pooser, Quantum sensing with squeezed light, *ACS Photonics* **6**, 1307 (2019).
- [32] Manuel Gessner, Augusto Smerzi, and Luca Pezzè, Multi-parameter squeezing for optimal quantum enhancements in sensor networks, *Nat. Commun.* **11**, 3817 (2020).

- [33] Kunal Sharma, Barry C. Sanders, and Mark M. Wilde, Optimal tests for continuous-variable quantum teleportation and photodetectors, arXiv e-prints, eid [ArXiv:2012.02754](https://arxiv.org/abs/2012.02754) (2020).
- [34] K. Sharma, M. Cerezo, Z. Holmes, L. Cincio, A. Sornborger, and P. J. Coles, Reformulation of the no-free-lunch theorem for entangled data sets, arXiv preprint [ArXiv:2007.04900](https://arxiv.org/abs/2007.04900) (2020).
- [35] L. Mandel and E. Wolf, *Optical Coherence and Quantum Optics* (Cambridge University Press, Cambridge, 1995).
- [36] H. Yuen and J. Shapiro, Optical communication with two-photon coherent states—part III: Quantum measurements realizable with photoemissive detectors, *IEEE Trans. Inf. Theory* **26**, 78 (1980).
- [37] Here, as elsewhere in this paper, we work in units where $\hbar\omega = 1$, where ω is the mode frequency.
- [38] A. Perelomov, *Generalized Coherent States and Their Applications* (Springer-Verlag, Berlin, 1985).
- [39] Jarrod R. McClean, Sergio Boixo, Vadim N. Smelyanskiy, Ryan Babbush, and Hartmut Neven, Barren plateaus in quantum neural network training landscapes, *Nat. Comm.* **9**, 4812 (2018).
- [40] M. Cerezo, Akira Sone, Tyler Volkoff, Lukasz Cincio, and Patrick J. Coles, Cost function dependent barren plateaus in shallow parametrized quantum circuits, *Nat. Commun.* **12**, 1791 (2021).
- [41] Alexey Uvarov and Jacob Biamonte, On barren plateaus and cost function locality in variational quantum algorithms, *J. Phys. A Math. Theor.* **54**, 245301 (2021).
- [42] Samson Wang, Enrico Fontana, M. Cerezo, Kunal Sharma, Akira Sone, Lukasz Cincio, and Patrick J. Coles, Noise-induced barren plateaus in variational quantum algorithms, arXiv preprint [ArXiv:2007.14384](https://arxiv.org/abs/2007.14384) (2020).
- [43] M. Cerezo and Patrick J. Coles, Higher order derivatives of quantum neural networks with barren plateaus, *Quantum Sci. Technol.* **6**, 035006 (2021).
- [44] Arthur Pesah, M. Cerezo, Samson Wang, Tyler Volkoff, Andrew T. Sornborger, and Patrick J. Coles, Absence of Barren Plateaus in Quantum Convolutional Neural Networks, *Phys. Rev. X* **11**, 041011 (2021).
- [45] Zoë Holmes, Andrew Arrasmith, Bin Yan, Patrick J. Coles, Andreas Albrecht, and Andrew T. Sornborger, Barren Plateaus Preclude Learning Scramblers, *Phys. Rev. Lett.* **126**, 190501 (2021).
- [46] Andrew Arrasmith, M. Cerezo, Piotr Czarnik, Lukasz Cincio, and Patrick J. Coles, *Quantum* **5**, 558 (2021).
- [47] Tyler Volkoff and Patrick J. Coles, Large gradients via correlation in random parameterized quantum circuits, *Quantum Sci. Technol.* **6**, 025008 (2021).
- [48] Carlos Ortiz Marrero, Mária Kieferová, and Nathan Wiebe, Entanglement induced barren plateaus, arXiv preprint [ArXiv:2010.15968](https://arxiv.org/abs/2010.15968) (2020).
- [49] Taylor L. Patti, Khadijeh Najafi, Xun Gao, and Susanne F. Yelin, Entanglement devised barren plateau mitigation, *Phys. Rev. Research* **3**, 033090 (2021).
- [50] Zoë Holmes, Kunal Sharma, M. Cerezo, and Patrick J. Coles, Connecting ansatz expressibility to gradient magnitudes and barren plateaus, arXiv preprint [ArXiv:2101.02138](https://arxiv.org/abs/2101.02138) (2021).
- [51] Edward Grant, Leonard Wossnig, Mateusz Ostaszewski, and Marcello Benedetti, An initialization strategy for addressing barren plateaus in parametrized quantum circuits, *Quantum* **3**, 214 (2019).
- [52] T. J. Volkoff, Efficient trainability of linear optical modules in quantum optical neural networks, *J. Russ. Laser Res.* **42**, 250 (2021).
- [53] To numerically compute the cost functions required for our simulations, we work in a truncated Hilbert space. Specifically, we truncate the Hilbert space to the 50 lowest lying Fock states.
- [54] Juan Miguel Arrazola, Thomas R. Bromley, Josh Izaac, Casey R. Myers, Kamil Brádler, and Nathan Killoran, Machine learning method for state preparation and gate synthesis on photonic quantum computers, *Quantum Sci. Technol.* **4**, 024004 (2019).
- [55] Nathan Killoran, Thomas R. Bromley, Juan Miguel Arrazola, Maria Schuld, Nicolás Quesada, and Seth Lloyd, Continuous-variable quantum neural networks, *Phys. Rev. Res.* **1**, 033063 (2019).
- [56] D. H. Wolpert and W. G. Macready, No free lunch theorems for optimization, *IEEE Trans. Evolutionary Comput.* **1**, 67 (1997).
- [57] D. H. Wolpert and W. G. Macready, No free lunch theorems for search, Tech. Rep. (Technical Report SFI-TR-95-02-010, Santa Fe Institute) (1995).
- [58] D. H. Wolpert, What the no free lunch theorems really mean: How to improve search algorithms, *Ubiquity* **2013**, 1 (2013).
- [59] D. H. Wolpert, The lack of a priori distinctions between learning algorithms, *Neural Comput.* **8**, 1341 (1996).
- [60] M. M. Wolf, *Mathematical foundations of supervised learning*, (2018).
- [61] K. Poland, K. Beer, and T. J. Osborne, No free lunch for quantum machine learning, arXiv preprint [ArXiv:2003.14103](https://arxiv.org/abs/2003.14103) (2020).
- [62] Andreas Winter and Dong Yang, Operational Resource Theory of Coherence, *Phys. Rev. Lett.* **116**, 120404 (2016).
- [63] L. Lami, Quantum data hiding with continuous variable systems, arXiv preprint [ArXiv:2102.01100](https://arxiv.org/abs/2102.01100) (2021).
- [64] K. K. Sabapathy and A. Winter, Bosonic data hiding: power of linear vs non-linear optics, arXiv preprint [ArXiv:2102.01622](https://arxiv.org/abs/2102.01622) (2021).
- [65] Nicolas C. Menicucci, Peter van Loock, Mile Gu, Christian Weedbrook, Timothy C. Ralph, and Michael A. Nielsen, Universal Quantum Computation with Continuous-Variable Cluster States, *Phys. Rev. Lett.* **97**, 110501 (2006).
- [66] D. Petz, *Quantum Information Theory and Quantum Statistics* (Springer, Berlin, 2008).
- [67] Andrew Lesniewski and Mary Beth Ruskai, Monotone riemannian metrics and relative entropy on noncommutative probability spaces, *J. Math. Phys.* **40**, 5702 (1999).
- [68] A. Serafini, *Quantum Continuous Variables* (CRC Press, Florida, 2017).
- [69] D. Braun, Invariant integration over the orthogonal group, *J. Phys. A: Math. Gen.* **39**, 14581 (2006).