

Theory for Equivariant Quantum Neural Networks

Quynh T. Nguyen,^{1,2,*} Louis Schatzki,^{3,4} Paolo Braccia⁵,^{1,5} Michael Ragone,^{1,6} Patrick J. Coles,¹ Frédéric Sauvage,¹ Martín Larocca,^{1,7} and M. Cerezo³

¹*Theoretical Division, Los Alamos National Laboratory, Los Alamos, New Mexico 87545, USA*

²*School of Engineering and Applied Sciences, Harvard University, Cambridge, Massachusetts 02138, USA*

³*Information Sciences, Los Alamos National Laboratory, Los Alamos, New Mexico 87545, USA*

⁴*Department of Electrical and Computer Engineering, University of Illinois at Urbana-Champaign, Urbana, Illinois 61801, USA*

⁵*Dipartimento di Fisica e Astronomia, Università di Firenze, Sesto Fiorentino, Florence 50019, Italy*

⁶*Department of Mathematics, University of California Davis, Davis, California 95616, USA*

⁷*Center for Nonlinear Studies, Los Alamos National Laboratory, Los Alamos, New Mexico 87545, USA*



(Received 5 February 2023; revised 10 March 2024; accepted 11 April 2024; published 6 May 2024)

Quantum neural network architectures that have little to no inductive biases are known to face trainability and generalization issues. Inspired by a similar problem, recent breakthroughs in machine learning address this challenge by creating models encoding the symmetries of the learning task. This is materialized through the usage of equivariant neural networks the action of which commutes with that of the symmetry. In this work, we import these ideas to the quantum realm by presenting a comprehensive theoretical framework to design equivariant quantum neural networks (EQNNs) for essentially any relevant symmetry group. We develop multiple methods to construct equivariant layers for EQNNs and analyze their advantages and drawbacks. Our methods can find unitary or general equivariant quantum channels efficiently even when the symmetry group is exponentially large or continuous. As a special implementation, we show how standard quantum convolutional neural networks (QCNNs) can be generalized to group-equivariant QCNNs where both the convolution and pooling layers are equivariant to the symmetry group. We then numerically demonstrate the effectiveness of a $SU(2)$ -equivariant QCNN over symmetry-agnostic QCNN on a classification task of phases of matter in the bond-alternating Heisenberg model. Our framework can be readily applied to virtually all areas of quantum machine learning. Lastly, we discuss about how symmetry-informed models such as EQNNs provide hopes to alleviate central challenges such as barren plateaus, poor local minima, and sample complexity.

DOI: [10.1103/PRXQuantum.5.020328](https://doi.org/10.1103/PRXQuantum.5.020328)

I. INTRODUCTION

Recognizing the underlying symmetries in a given data set has played a fundamental role in classical machine learning. For instance, noting that the picture of a cat still depicts a cat when we translate the pixels of the image gives a hint as to why convolutional neural networks [1] have been so successful in image classification: they process images in a translationally symmetric way [2].

In recent years, the importance of symmetries in machine learning has been studied in problems with more

general symmetry groups than translations, leading to the burgeoning field of geometric deep learning [2]. The central thesis of this field is that prior symmetry knowledge should be incorporated into the model, thus effectively constraining the search space and easing the learning task. Indeed, symmetry-respecting models have been observed to perform and generalize better than problem-agnostic ones in a wide variety of tasks [2–9]. As such, a great deal of work has also gone into developing a mathematically rigorous framework for designing symmetry-informed models through the machinery of representation theory. This has provided the basis for so-called *equivariant neural networks* (ENNs) [10–14], the key property of which is that their action commutes with that of the symmetry group. In other words, applying a symmetry transformation to the input and then sending it through the ENN produces the same result as sending the raw input through the ENN and then applying the transformation.

*Corresponding author: qnguyen@g.harvard.edu

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.

Recently, some of the ideas of geometric deep learning have been imported to the field of *quantum machine learning* (QML) [15–19]. QML has become a rapidly growing framework to make practical use of noisy intermediate-scale quantum devices [20]. Here, the hope is that by accessing the exponentially large Hilbert space, quantum models can obtain a computational advantage over their classical counterparts [21,22], especially for quantum data [23–25]. Despite its promise, there are still several challenges that one needs to address before unlocking the full potential of QML. In particular, a growing amount of evidence suggests that models with little to no inductive biases have poor trainability and generalization, greatly limiting their scalability [26–41].

Geometric quantum machine learning (GQML) attempts to solve the aforementioned issues by leveraging ideas from geometric deep learning to construct quantum models with sharp inductive biases based on the symmetries of the problem at hand. For instance, when classifying between states presenting a large, or a low, amount of multipartite entanglement [24,42,43], it is natural to employ models the outputs of which remain invariant under the action of any local unitary [44]. While recent proposals have started to reveal the power of GQML [44–52], the field is still in its infancy and a more systematic approach to symmetry-encoded model design is needed.

The goal of this work is to offer a theoretical framework for building GQML models based on extending the notion of classical ENNs to *equivariant quantum neural networks* (EQNNs) (see Fig. 1). Our main contributions can be summarized as follows:

- (a) We provide an interpretation for EQNN layers as a form of generalized Fourier-space action, meaning that they perform a group Fourier transform, act on

the Fourier components, and transform back. This allows us to quantify the number of free parameters in an EQNN layer and it unravels the exciting possibility of using different group representations as hyperparameters to act on different generalized Fourier spaces (Sec. IV).

- (b) We introduce a general framework for EQNNs, extending previous results from unitary quantum neural networks to channels. We characterize the different types of EQNN layers such as standard, embedding, pooling, lifting and projection layers. This permits a classification of EQNN layers depending on their input and output representations. In addition, we also explore how equivariant nonlinearities can be introduced via multiple copies of the data (Sec. IV).
- (c) We describe three alternative methods for constructing and parametrizing EQNNs. These are based on finding the *null space* of a certain system of matrix equations, applying the *twirling* formula over the symmetry group and using the *Choi-operator* representation of channels. Our methods have a better complexity than existing methods and can efficiently find unitary or nonunitary equivariant layers even when the symmetry group is exponentially large. We discuss strengths and weaknesses of each approach and general methods to optimize and/or train equivariant channels (Sec. V).
- (d) We exemplify our techniques by showing how to generalize standard quantum convolutional neural networks (QCNNs) to group-equivariant QCNNs where the convolutional and pooling layers are equivariant to the symmetry group of the task. As examples, we show how to construct $\mathbb{Z}_2 \times \mathbb{Z}_2$ and \mathbb{Z}_n -equivariant EQNNs (Sec. VI). Moreover,

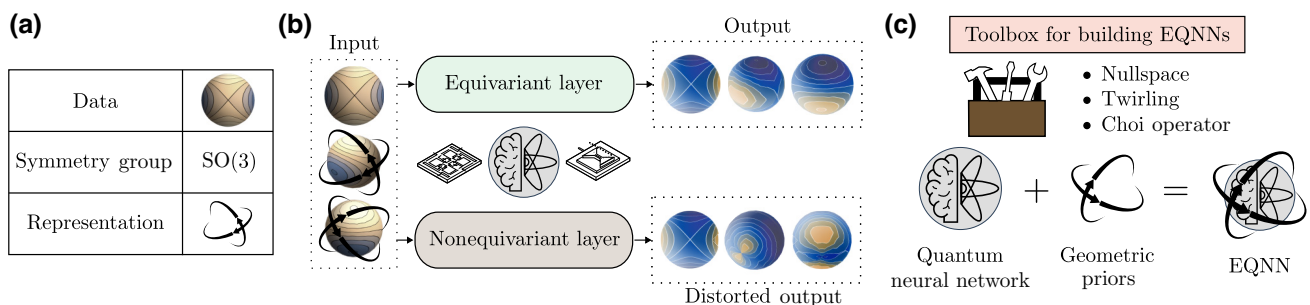


FIG. 1. A schematic representation of our main results. (a) In GQML, we start by identifying the symmetry group—or groups—that leave the data labels invariant. For the example shown, the data can be visualized on a three-dimensional sphere and the labels are invariant under the action of $\mathbb{S}\mathbb{O}(3)$. (b) Both in classical and quantum machine learning, it has been shown that models with equivariant layers often have an improved performance over nonequivariant architectures. The key feature of equivariance is that applying a rotation to the input data and sending it through the layer is the same as first sending the data through the layer and then rotating the output. On the other hand, feeding either a raw or a rotated data instance into a nonequivariant layer usually leads to distorted outputs that are not related by a rotation. (c) In this work, we provide a toolbox of methods for creating equivariant quantum neural networks (EQNNs) that can be readily used to construct quantum architectures with strong geometric priors.

we present a new architecture called $\mathbb{S}\mathbb{U}(2)$ -equivariant QCNN and numerically demonstrate its advantage over symmetry-agnostic QCNNs in a quantum phase-classification task of the bond-alternating Heisenberg model on up to 13 qubits (Sec. VII).

We conclude with a discussion on how EQNNs provide hope to alleviate critical challenges in QML such as ill-behaved training landscapes (barren plateaus and local minima) and to reduce the sample complexity (data requirements) of the model. Taken together, we hope that our results will serve as blueprints and guidelines for a more representation-theoretic approach to QML.

II. RELATED WORK

A. Equivariance in geometric deep learning

In this section, we provide an overview of the literature on classical equivariant neural networks (ENNs), leaving the formal treatment of equivariance to Sec. III.

At a high level, equivariance is a mathematical property that preserves symmetries in features throughout a multilayer ENN. One imposes equivariance onto ENN layers via tools from group representation theory, the workhorse behind geometric deep learning [2]. The most well-known equivariant architecture is the convolutional neural network (CNN) [1], ubiquitous in image and signal processing. The relevant symmetry group in CNNs is the translation group in the plane \mathbb{R}^2 and one can show that their convolution and pooling layers are equivariant to this group [11]. Ideas to generalize CNNs to other groups and data have first been laid out in Ref. [10] and have further been made mathematically rigorous in Refs. [11,12]. These works are concerned with the so-called homogeneous ENNs, which include, as special cases, spherical CNNs (where the relevant group is $\mathbb{S}\mathbb{O}(3)$) for spherical images [6] and Euclidean neural networks (the Euclidean group $\mathbb{E}(n) = \mathbb{R}^n \rtimes \mathbb{O}(n)$ and its subgroups) for molecular data [53–56]. Notable nonhomogeneous architectures include graph neural networks (the permutation group S_n) [4,57,58]. In addition, more advanced representation-theoretic treatments on nonhomogeneous data [14] have led to steerable CNNs [59] and gauge-equivariant CNNs [13] on general manifolds. Moreover, it has been shown that equivariant layers can be constructed from either the group-space or Fourier-space perspectives [5]. More recently, methods for designing equivariant layers have been studied in Refs. [8,51,60,61]. For a theoretical analysis of the improvements in training and generalization error arising from using ENNs, we refer the reader to Refs. [9,62–64], while the expressibility and universality of ENNs have been studied in Refs. [65–69].

B. Equivariance in quantum information

Equivariance [70] has a long history in quantum information theory [71–91]. As such, we will not attempt here to review its full impact on the field but, rather, we will focus on several relevant works where equivariance has been studied in the context of quantum channels.

To begin, the set of all irreducible $\mathbb{S}\mathbb{U}(2)$ -equivariant channels has been characterized in Ref. [92], with extensions to a wide class of finite groups presented in Ref. [93]. The work in Ref. [94] presents conditions to construct group-equivariant generalized-extreme channels. On the other hand, the history of equivariance in QML is much more recent. In Refs. [44,45], the authors lay a theoretical groundwork for the integration of symmetries into QML. However, prior works are either nonconstructive [44] or only work efficiently on restricted sets of problems and symmetries [45]. In particular, two main types of symmetries have been most extensively explored: the action of the local unitary group $\mathbb{S}\mathbb{U}(d)$ on each qudit in a correlated manner $U^{\otimes n}$ and the action of the permutation group S_n by permuting the qudits. It is well-known fact in representation theory, called the Schur-Weyl duality, that these two group representations commute.

On the side of $\mathbb{S}\mathbb{U}(d)$ symmetry, in Ref. [49], the authors have proposed a specific task—approximating matrix elements of S_n irreps evaluated on arbitrary group-algebra elements—together with a novel polynomial-time quantum algorithm, based on the combination of the quantum Schur transform and Hamiltonian simulation, that potentially achieves an exponential quantum speed-up given that the best known classical algorithms require $\mathcal{O}(n!n^2)$ time. In turn, Ref. [48] exploits the ideas in Ref. [49] to derive an ansatz—the S_n -equivariant convolutional quantum alternating ansatz (S_n CQA)—that is universal for the subgroup of $\mathbb{S}\mathbb{U}(d)$ -equivariant unitaries. As the name suggests, it is based on the qudit permutation action of S_n on the quantum system that, via Schur-Weyl duality, linearly spans the subspace of $\mathbb{S}\mathbb{U}(d)$ -equivariant operators. Interestingly, S_n CQA can be shown to achieve universality in the subgroup of symmetric unitaries with only four-body interactions, something that is remarkable given the typical limitations of universality imposed by locality constraints [88,95]. The performance and resource requirements of S_n CQA are benchmarked in Ref. [50].

On the S_n -symmetry side, Ref. [96] has shown that S_n -equivariant QNNs exhibit a wide range of favorable properties, such as being immune to barren plateaus, efficiently reaching overparametrization, and being able to generalize well from few training points. Moreover, methods for constructing equivariant quantum circuits for graph problems have been given in [46,47,52,97]. We also note that, while not explicitly mentioned, some recent quantum algorithms can be analyzed from an equivariance point of view [21,23,98–100]. We discuss these in Appendix A.

III. PRELIMINARIES

Here, we give some of the necessary background in QML and representation theory to tackle GQML. For a more comprehensive treatment of these topics, we refer the reader to the standard textbooks in representation theory [101–103] and the geometric ML-theory literature [2,12,14]. For a QML-oriented approach to group theory and representation theory, see Ref. [104].

A. From QML to GQML

For simplicity and concreteness, in this paper we focus on quantum supervised learning with scalar labels. However, we remark that GQML is relevant in other contexts, such as unsupervised learning [105,106], generative modeling [107–110], or reinforcement learning [111,112]. For instance, the constructions presented here can be readily adapted to learning problems with nonscalar output (e.g., quantum generative models, where the output is a quantum state or a probability distribution).

Suppose that we are given some data set composed of M quantum states and scalar labels $\{\rho_i, y_i\}_{i=1}^M$, where $\rho_i \in \mathcal{R}$ are quantum states from a data domain $\mathcal{R} \subset \mathcal{B}(\mathcal{H})$, with $\mathcal{B}(\mathcal{H})$ the set of bounded linear operators on \mathcal{H} . The labels come from a label domain \mathcal{Y} and are obtained from a (potentially probabilistic) function:

$$f : \mathcal{R} \rightarrow \mathcal{Y} \quad (\text{Underlying function}). \quad (1)$$

For example, in binary classification one has $\mathcal{Y} = \{0, 1\}$. These data may come from some physical quantum mechanical process (quantum data [24]) or may be classical information embedded into quantum states (classical data [113]). Given the data set, one then optimizes a learning model:

$$h_\theta : \mathcal{R} \rightarrow \mathcal{Y} \quad (\text{Model}), \quad (2)$$

where θ are trainable parameters, with the intent of closely approximating the underlying function f .

In variational QML [18], the states in the data set are fed into a trainable quantum circuit, which is usually modeled by a sequence of parametrized unitary matrices. However, in this work, we will consider more general operations—parametrized quantum channels—which we refer to as *quantum neural networks* (QNNs). Respectively denoting the spaces of bounded linear operators in \mathcal{H}^{in} and \mathcal{H}^{out} as $\mathcal{B}^{\text{in}} := \mathcal{B}(\mathcal{H}^{\text{in}})$ and $\mathcal{B}^{\text{out}} := \mathcal{B}(\mathcal{H}^{\text{out}})$, a QNN is a parametrized completely positive and trace-preserving (CPTP) linear map $\mathcal{N}_\theta : \mathcal{B}^{\text{in}} \rightarrow \mathcal{B}^{\text{out}}$.

We can further decompose the QNN as a concatenation of channels or *layers*. We say that \mathcal{N}_θ is an L -layered QNN if it can be expressed as $\mathcal{N}_\theta = \mathcal{N}_{\theta_L}^L \circ \dots \circ \mathcal{N}_{\theta_1}^1$, where the $\mathcal{N}_{\theta_l}^l$ (with $l = 1, \dots, L$) are parametrized CPTP channels such that $\theta = (\theta_1, \dots, \theta_L)$. From the foregoing, the

l th layer maps between operators acting on some Hilbert space \mathcal{H}^{l-1} to operators acting on some (potential different) Hilbert space \mathcal{H}^l . That is, $\mathcal{N}_{\theta_l}^l : \mathcal{B}^{l-1} \rightarrow \mathcal{B}^l$, where, for simplicity of notation, we have defined $\mathcal{B}^l := \mathcal{B}(\mathcal{H}^l)$.

After applying the QNN to an input state ρ , one measures the resulting state with respect to a set of observables $\{O_j\}_j$ to obtain the expectation values $\{\text{Tr}[\mathcal{N}_\theta(\rho)O_j]\}_j$. Finally, a classical postprocessing step, \mathcal{C} , maps these outcomes to a loss function

$$\ell_\theta(\rho) = \mathcal{C}(\{\text{Tr}[\mathcal{N}_\theta(\rho)O_j]\}_j). \quad (3)$$

We quantify the performance of the model over the data set via the so-called empirical loss,

$$\widehat{\mathcal{L}}_\theta(\{\rho_i, y_i\}_{i=1}^M) = \frac{1}{M} \sum_{i=1}^M \mathcal{F}(\ell_\theta(\rho_i), y_i), \quad (4)$$

defined in terms of some problem-dependent function \mathcal{F} . Finally, employing a classical computer, one optimizes over the parameters θ to minimize the empirical loss until certain convergence conditions are met. The optimal parameters, along with the loss function, are used to predict labels.

One of the most important aspects that make or break the QML scheme are its inductive biases, i.e., the assumptions about the problem that one embeds in the structure of the model. In our case, this amounts to an adequate choice of the parametrized layers $\mathcal{N}_{\theta_l}^l$ forming the QNN and of the measurement operators O_j . In a nutshell, the inductive biases are responsible for the model exploring only a subset of all possible functions from \mathcal{R} to \mathcal{Y} . If these inductive biases are too general or not accurate, the model is expected to train poorly, while models with appropriate inductive biases can often benefit from an improved performance [26,31,96]. GQML aims at providing a framework for incorporating prior geometrical information in the model with the hope of improving its trainability, data requirements, generalization, and overall performance. In particular, the main goal of GQML is to create models respecting the underlying symmetries of the domain over which they act. In the next sections, we will briefly review how to use tools from representation theory to deal with symmetries, as well as recall basic concepts such as equivariance and invariance.

B. Symmetry groups and representation theory

The first step toward building a GQML model is to identify the set of relevant operations that the model needs to preserve. We say that a QML problem has symmetry with respect to a group G if the labels are unchanged under the action of a representation of G on the input states.

Definition 1 (Label symmetries and G -invariance). Given a compact group G and some unitary representation

R acting on quantum states ρ , we say that the underlying function f has a label symmetry if it is G -invariant, i.e., if

$$f(R(g)\rho R(g)^\dagger) = f(\rho), \quad \forall g \in G. \quad (5)$$

As previously mentioned, the goal of GQML is to build models that respect the label symmetries of the data. That is, we want to build G -invariant models such that $h_\theta(\rho) = h_\theta(R(g)\rho R(g)^\dagger)$, for any $g \in G$, and for all values of θ .

To further understand how symmetry groups act and how one can manipulate them, we recall here some basic concepts from representation theory (for further background, see Ref. [104]). Namely, given a group G , its *representation* describes its action on some vector space \mathcal{H} , which we assume for simplicity to be a Hilbert space.

Definition 2 (Representation). A representation (R, \mathcal{H}) of a group G on a vector space \mathcal{H} is a homomorphism $R: G \rightarrow GL(\mathcal{H})$ from the group G to the space of invertible linear operators on \mathcal{H} that preserves the group structure of G .

Specifically, a group homomorphism R satisfies

$$R(g_1)R(g_2) = R(g_1g_2) \quad \forall g_1, g_2 \in G. \quad (6)$$

This implies that, for all $g \in G$, the representation of its inverse is the inverse of its representation, $R(g^{-1}) = R(g)^{-1}$, and the representation of the identity element e is the identity operator on \mathcal{H} , $R(e) = \mathbb{1}_{\dim(\mathcal{H})}$. Given a representation, it is relevant to define its commutant.

Definition 3 (Commutant). Given a representation R of G , we define the commutant of R as the set of bounded linear operators on \mathcal{H} that commute with every element in R , i.e.,

$$\text{comm}(R) = \{H \in \mathcal{B}(\mathcal{H}) \mid [H, R(g)] = 0 \forall g \in G\}. \quad (7)$$

Consider the following remarks about representations:

- A representation is *faithful* if it maps distinct group elements to distinct elements in \mathcal{H} . As an example of unfaithfulness, the *trivial representation* maps all group elements to the identity in \mathcal{H} .
- Two representations, R_1 and R_2 , are *equivalent* if there exists a change of basis W such that $WR_1(g)W^\dagger = R_2(g)$ for all $g \in G$, in which case we denote $R_1 \cong R_2$.
- A *subrepresentation* is a subspace $\mathcal{K} \subset \mathcal{H}$ that is invariant under the action of the representation, i.e., $R(g)|w\rangle \in \mathcal{K}$ for all $g \in G$ and $|w\rangle \in \mathcal{K}$. The group can then be represented through $R|_{\mathcal{K}}$, the restriction

of R to the vector subspace \mathcal{K} . A subrepresentation \mathcal{K} is nontrivial if $\mathcal{K} \neq \{0\}$ (the zero vector) and $\mathcal{K} \neq \mathcal{H}$.

Definition 4. (Irreps) A representation is said to be an irreducible representation (irrep) if it contains no nontrivial subrepresentations.

Irreps are the fundamental building blocks of representation theory. For any finite or compact group, the representations can be chosen to be unitary [114]. Hence in the rest of this paper, we will consider unitary representations on complex Hilbert spaces. In the case that the representation is finite-dimensional, we can go a step further and say that the vector space can be decomposed into a direct sum over irreducible subrepresentations. This leads to the so-called *isotypic* decomposition,

$$\mathcal{H} \cong \bigoplus_{\lambda} \mathcal{H}_{\lambda} \otimes \mathbb{C}^{m_{\lambda}}, \quad R(g) \cong \bigoplus_{\lambda} R_{\lambda}(g) \otimes \mathbb{1}_{m_{\lambda}}, \quad (8)$$

where \cong indicates that there exists a global change of basis matrix W that simultaneously block diagonalizes the unitaries $R(g)$ for all $g \in G$. Here, λ labels the irreps, m_{λ} is the multiplicity of the irrep R_{λ} , and $R_{\lambda}(g) \in \mathbb{C}^{d_{\lambda} \times d_{\lambda}}$. Note that $\sum_{\lambda} d_{\lambda} m_{\lambda} = \dim(\mathcal{H})$.

When G is not a finite group, we assume it to be a compact Lie group with an associated Lie algebra \mathfrak{g} such that $e^{\mathfrak{g}} = G$. That is, $\mathfrak{g} = \{a|e^a \in G\}$. In particular, if G has a representation R , then \mathfrak{g} has a representation r given by the differential of R , i.e., given $a \in \mathfrak{g}$, $R(e^a) = e^{r(a)}$.

In this work, we will mainly focus on the *adjoint representation* of G , as it describes how the group acts on density matrices (and other bounded operators). A unitary representation R on \mathcal{H} induces an action on $\mathcal{B}(\mathcal{H})$, given by

$$\text{Ad}_{R(g)}(\rho) = R(g)\rho R(g)^\dagger, \quad \forall g \in G, \rho \in \mathcal{B}(\mathcal{H}). \quad (9)$$

where $\text{Ad}_{R(g)}$ denotes the adjoint representation. Note that for the case of Lie groups, the adjoint representation also exists at the Lie-algebra level and is given by $\text{ad}_{r(a)}(\cdot) = [r(a), \cdot]$.

To finish this section, we find it convenient to define a distinction between symmetry groups.

Definition 5 (Inner and outer symmetries). Given a composite Hilbert space, we call a representation of a group an *inner* symmetry if it acts locally on each subsystem and an *outer* symmetry if it permutes the subsystems.

For instance, when working with n -qubit systems, the tensor representation of $\text{SU}(2)$, $R(g \in \text{SU}(2)) = g^{\otimes n}$ is an inner symmetry, as it acts locally on each subsystem. On the other hand, the qubit-permuting representation of S_n , given by $R(g) \otimes_{j=1}^n |\psi_j\rangle = \otimes_{j=1}^n |\psi_{g^{-1}(j)}\rangle$ is an outer symmetry.

C. Equivariance and invariance in quantum neural networks

Here, we present a recipe to obtain G -invariant QML models based on the key concept of *equivariance*, which we will first define for linear maps and then for operators.

Given a group G and a representations R , we typically say that a linear map $\phi : \mathcal{B}(\mathcal{H}) \rightarrow \mathcal{B}(\mathcal{H})$ is equivariant if and only if $\phi \circ \text{Ad}_{R(g)} = \text{Ad}_{R(g)} \circ \phi$ for all $g \in G$. We can extend this definition by noting that neither the input and output representations, nor the Hilbert spaces, need to be the same. Thus, we consider the following general definition.

Definition 6 (Equivariant map). Given a group G and its representations $(R^{\text{in}}, \mathcal{H}^{\text{in}})$ and $(R^{\text{out}}, \mathcal{H}^{\text{out}})$, a linear map $\phi : \mathcal{B}^{\text{in}} \rightarrow \mathcal{B}^{\text{out}}$ is $(G, R^{\text{in}}, R^{\text{out}})$ -equivariant if and only if

$$\phi \circ \text{Ad}_{R^{\text{in}}(g)} = \text{Ad}_{R^{\text{out}}(g)} \circ \phi, \quad \forall g \in G. \quad (10)$$

The property of equivariance can be visualized via the following commutative diagram:

$$\begin{array}{ccc} \mathcal{B}^{\text{in}} & \xrightarrow{\text{Ad}_{R^{\text{in}}(g)}} & \mathcal{B}^{\text{in}} \\ \downarrow \phi & & \downarrow \phi \\ \mathcal{B}^{\text{out}} & \xrightarrow{\text{Ad}_{R^{\text{out}}(g)}} & \mathcal{B}^{\text{out}} \end{array} .$$

The action of an equivariant map ϕ *commutes* with the action of the group. That is, for an equivariant ϕ , it is equivalent to (i) first acting with ϕ and then acting with R^{out} or to (ii) first acting with R^{in} and then acting with ϕ . Note that for the special case of R^{out} being the trivial representation, the map is *invariant*, such that $\phi \circ \text{Ad}_{R^{\text{in}}(g)} = \phi$ for all $g \in G$.

Next, let us define what an equivariant operator is.

Definition 7 (Equivariant operator). Given a group G and its representation (R, \mathcal{H}) , an operator $O \in \mathcal{B}(\mathcal{H})$ is (G, R) -equivariant if and only if

$$[O, R(g)] = 0, \quad \forall g \in G. \quad (11)$$

Evidently, Definition 7 implies that $O \in \text{comm}(R)$, from which we can easily see that $\text{comm}(R)$ is the space of all equivariant operators. Moreover, we can also see that the adjoint action of a (G, R) -equivariant operator is a (G, R, R) -equivariant map. That is, $\text{ad}_O \circ \text{Ad}_{R(g)} = \text{Ad}_{R(g)} \circ \text{ad}_O$.

The previous definitions present us with a recipe to build GQML models of the form in Eq. (3), the outputs of which are invariant under the action of the group.

Proposition 1 (Invariance from equivariance). A model consisting of an $(G, R^{\text{in}}, R^{\text{out}})$ -equivariant QNN and a (G, R^{out}) -equivariant set of measurements is G -invariant.

Proof. For every $g \in G$, $\rho \in \mathcal{B}^{\text{in}}$, and θ , we have

$$\begin{aligned} h_{\theta}(\text{Ad}_{R^{\text{in}}(g)}(\rho)) &= \mathcal{C}(\{\text{Tr}[\mathcal{N}_{\theta}(\text{Ad}_{R^{\text{in}}(g)}(\rho))O_j]\}_j) \\ &= \mathcal{C}(\{\text{Tr}[\text{Ad}_{R^{\text{out}}(g)}(\mathcal{N}_{\theta}(\rho))O_j]\}_j) \\ &= \mathcal{C}(\{\text{Tr}[\mathcal{N}_{\theta}(\rho)R^{\text{out}}(g)^{\dagger}O_jR^{\text{out}}(g)]\}_j) \\ &= \mathcal{C}(\{\text{Tr}[\mathcal{N}_{\theta}(\rho)O_j]\}_j) = h_{\theta}(\rho). \quad \blacksquare \end{aligned}$$

Armed with the previous definitions, we are now ready to present the basic framework for EQNNs.

First, however, we find it instructive to provide an example of a classification problem naturally amenable to these methods. Suppose that we are given the ground states of the bond-alternating Heisenberg model, which has the Hamiltonian

$$H = J_1 \sum_{i \text{ even}} S_i \cdot S_{i+1} + J_2 \sum_{i \text{ odd}} S_i \cdot S_{i+1}, \quad (12)$$

where $S_i = (S_x^i, S_y^i, S_z^i)$ is the spin operator for the i th qubit. There are two phases of matter for this Hamiltonian: trivial and topologically protected. As a learning problem, we consider the task of determining if the states are in the trivial or topologically protected phases. Consider the representation $R(g) = g^{\otimes n}$ of $\mathbb{S}\mathbb{U}(2)$. For a ground state $|\psi\rangle$, one can show that $R(g)|\psi\rangle$ is also a ground state. Thus, the labels of states are invariant under an action of $\mathbb{S}\mathbb{U}(2)$. In Sec. VII, we return to this problem and show that a EQNN significantly outperforms a quantum convolutional neural network for this task.

IV. THEORY OF EQUIVARIANT LAYERS FOR EQNNs

In this section, we will shed some light on the importance of the choice of representation by studying how EQNNs act on data and how many degrees of freedom they have. Most notably, we will show that layers that are equivariant to different representations can process data in different ways, so that a given layer could potentially “see” information that is inaccessible to another one. The latter will point to the crucial importance that intermediate representations have. Finally, we will present a classification of EQNN layers based on their input and output representations, allowing them to be nonlinear, and even change the symmetry group under which they are equivariant. Our results can be summarized in Fig. 2.

A. Equivariant layers as Fourier-space actions

Let us start by analyzing how EQNNs act on data. For simplicity, we first consider the case in which $R^{\text{in}} = R^{\text{out}} = R$, the states in the data set are pure $\rho = |\psi\rangle\langle\psi|$, and the EQNN is unitary, i.e., $\mathcal{N}_{\theta}(\rho) = U(\theta)\rho U(\theta)^{\dagger}$. Note that if \mathcal{N}_{θ} is a (G, R, R) -equivariant map, then $U(\theta)$ is a (G, R) -equivariant operator and, hence, it belongs to $\text{comm}(R)$.

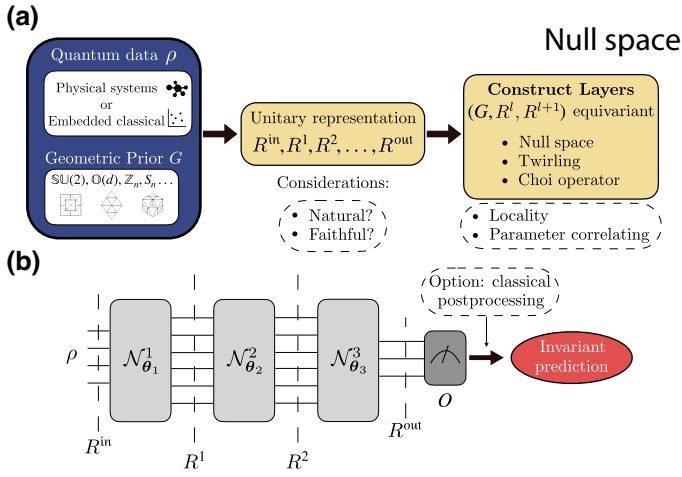


FIG. 2. The equivariant quantum neural network. (a) The design procedure. We consider a QML problem composed of a data set (that can either be quantum mechanical in nature or corresponding to classical data that have been encoded in quantum states) as well as a label symmetry group G . The first step is to define the input and output representations of G at each layer, where these can be natural, faithful, nonfaithful, etc. From here, we will provide different techniques that allow us to construct the EQNN layers and control, e.g., the locality of their gates. (b) The EQNN architecture. The dashed lines indicate the representations of the symmetry group G at specific stages in the EQNN, which may change between layers. At first, the input state ρ_{in} is acted upon by the representation R^{in} . The l th layer of the EQNN, $\mathcal{N}_{\theta^l}^l$, must be (G, R^l, R^{l+1}) -equivariant. In sum, the full architecture, $\phi = \mathcal{N}_{\theta^L}^L \circ \dots \circ \mathcal{N}_{\theta^1}^1$, is $(G, R^{\text{in}}, R^{\text{out}})$ -equivariant. The (G, R^{out}) -equivariant measurement operator O is in the commutant of the output representation R^{out} . Note that if we only want the EQNN to produce an output state equivariantly or invariantly (e.g., in generative models), we can omit the measurements.

Hence, we can understand the action of $U(\theta)$ on $|\psi\rangle$ by studying the structure of the commutant.

Theorem 1 (Structure of commutant (Theorem IX.11.2 in Ref. [101])). Let R be a unitary representation of a finite-dimensional compact group G . Then, under the same change of basis W , which block diagonalizes R as in Eq. (8), any operator $H \in \text{comm}(R)$ takes the following block-diagonal form:

$$H \cong \bigoplus_{\lambda} \mathbb{1}_{d_{\lambda}} \otimes H_{\lambda}, \quad (13)$$

where each H_{λ} is an m_{λ} -dimensional operator that is repeated d_{λ} times.

The previous theorem shows that any equivariant unitary can be expressed as $U(\theta) = W^{\dagger} \left(\bigoplus_{\lambda} \mathbb{1}_{d_{\lambda}} \otimes U_{\lambda}(\theta) \right) W$, indicating that in the irrep basis, it can only act nontrivially on the multiplicity space. Drawing a parallelism with the classical ML literature, where it has been shown that linear

equivariant maps can only act on the group Fourier components of the data [5,6,53,57,58], we can also here interpret EQNNs as a form of generalized Fourier-space action. Specifically, the action of $U(\theta)$ can be understood as (i) first transforming the data to the generalized Fourier space $W|\psi\rangle = \bigoplus_{\lambda} |\bar{\psi}_{\lambda}\rangle \otimes |\psi_{\lambda}\rangle$, (ii) acting on each Fourier component $|\psi_{\lambda}\rangle$ with $U_{\lambda}(\theta)$, and (iii) transforming back with W^{\dagger} . That is,

$$U(\theta) |\psi\rangle = W^{\dagger} \left(\bigoplus_{\lambda} |\bar{\psi}_{\lambda}\rangle \otimes U_{\lambda}(\theta) |\psi_{\lambda}\rangle \right). \quad (14)$$

Note that this interpretation can be readily generalized to channels.

Here, we can see that once the representation of G is fixed, so is the information in the input state to which one has access (equivariantly). Explicitly, the EQNN cannot manipulate information stored in the components $|\bar{\psi}_{\lambda}\rangle$ of the input state. As we will see in Sec. IV B, one can still try to access this information via changes of representation.

Notably, Eq. (14) generalizes *group convolution* in the Fourier basis: when R is the regular representation, the change of basis is the well-known group Fourier transform [91,115] (see Appendix B). This generalized Fourier-space picture has proved crucial in designing various classical architectures [5,6,53,57,58]. This also provides a representation-theoretic justification for the recent quantum “convolutional layers” in Ref. [23]. Recently, this interpretation of equivariant unitaries has also been noted for the special case of $\text{SU}(d)$ -equivariant quantum circuits in Refs. [48,49].

B. Free parameters in EQNNs

1. Equivariant unitaries

The Fourier-space picture previously discussed enables the counting of free parameters in equivariant unitaries.

Theorem 2 (Free parameters in equivariant unitaries). Under the same setup as Theorem 1, the unitary operators in $\text{comm}(R)$ can be fully parametrized by $\sum_{\lambda} m_{\lambda}^2$ real scalars.

Proof. Any unitary U in $\text{comm}(R)$ takes the block-diagonal form $U = \bigoplus_{\lambda} \mathbb{1}_{d_{\lambda}} \otimes U_{\lambda}$ in the Fourier basis. Observe that the operators U_{λ} must also be unitaries, since $U^{\dagger}U = \bigoplus_{\lambda} \mathbb{1}_{d_{\lambda}} \otimes U_{\lambda}^{\dagger}U_{\lambda}$. A unitary in $\mathbb{U}(m_{\lambda})$ is parametrized by m_{λ}^2 real scalars; hence, a total number of $\sum_{\lambda} m_{\lambda}^2$ parameters suffice to parametrize U . ■

Theorem 2 describes how “significant” the symmetry is to the problem, in the sense that the larger the representations of G , the smaller is the commutant and thus the fewer parameters are needed to fully characterize equivariant unitaries. In Table I, we present examples of different

TABLE I. Free parameters in unitary EQNNs. We show how different symmetries impact the number of free parameters in a (G, R, R) -equivariant unitary. Here, for set a S , we have defined $\mathbb{C}[S] \equiv \text{span}_{\mathbb{C}}(S)$.

Group	Representation	Free parameters	$\text{comm}(R)$
None	$R_{\text{trivial}}(g) = \mathbb{1}_{2^n}$	4^n	$\mathbb{C}[\mathbb{U}(2^n)]$
$\mathbb{U}(2^n)$	$R_{\text{def}}(g) = g$	1	$\mathbb{C}[\mathbb{1}_{2^n}]$
$\mathbb{U}(2)$	$R_{\text{tens}}(g) = g^{\otimes n}$	$\frac{1}{n+2} \binom{2n+2}{n+1} \in \Omega(2^n)$	$\mathbb{C}[R_{\text{qub}}(S_n)]$
S_n	$R_{\text{qub}}(g) \otimes_{i=1}^n \psi_i\rangle = \otimes_{i=1}^n \psi_{g^{-1}(i)}\rangle$	$\binom{n+3}{3} \in \Theta(n^3)$	$\mathbb{C}[R_{\text{tens}}(\mathbb{U}(2))]$

symmetries constraining the number of free parameters in a unitary EQNN to both exponentially many, polynomially many, and constant.

2. Equivariant channels

We have already seen how the inductive biases in unitary EQNNs affect their structure and, concomitantly, their number of free parameters. We now turn our attention to $(G, R^{\text{in}}, R^{\text{out}})$ -equivariant channels. First, recall that any linear channel $\phi : \mathcal{B}^{\text{in}} \rightarrow \mathcal{B}^{\text{out}}$, can be fully characterized through its Choi operator [116]

$$J^\phi = \sum_{i,j} |i\rangle\langle j| \otimes \phi(|i\rangle\langle j|), \quad (15)$$

which acts in $\mathcal{B}(\mathcal{H}^{\text{in}} \otimes \mathcal{H}^{\text{out}})$. The action of ϕ on an input state $\rho \in \mathcal{B}^{\text{in}}$ can be recovered from J^ϕ as follows: [117]

$$\phi(\rho) = \text{Tr}_{\text{in}}[J^\phi(\rho^\top \otimes \mathbb{1}_{\dim(\mathcal{H}^{\text{out}})})]. \quad (16)$$

The Choi operator is related to equivariance via the following theorem.

Lemma 1 (Lemma 11 in Ref. [78] paraphrased). A channel ϕ is $(G, R^{\text{in}}, R^{\text{out}})$ -equivariant if and only if $J^\phi \in \text{comm}(R^{\text{in}*} \otimes R^{\text{out}})$, where the asterisk (*) denotes the complex conjugate.

Noting that $(R^{\text{in}*} \otimes R^{\text{out}})$ is a valid representation as per Definition 2, we can combine Theorem 1 and Lemma 1 to determine a parameter count for general equivariant channels.

Theorem 3 (Free parameters in equivariant channels). Let the irrep decomposition of $R := R^{\text{in}*} \otimes R^{\text{out}}$ be $R(g) \cong \bigoplus_q R_q(g) \otimes \mathbb{1}_{m_q}$. Then, any $(G, R^{\text{in}}, R^{\text{out}})$ -equivariant CPTP channels can be fully parametrized via $\sum_q m_q^2 - C(R^{\text{in}}, R^{\text{out}})$ real scalars, where $C(R^{\text{in}}, R^{\text{out}})$ is a positive constant that depends on the considered representations.

We defer the proof to Appendix C. Intuitively, the parameter count of equivariant CP maps follows similarly

to the proof of Theorem 2 and the extra term $C(R^{\text{in}}, R^{\text{out}})$ arises from imposing that the channel ϕ must be trace preserving (TP).

Similar to the classical ML literature [59], the parameter-count benefit of using equivariant layers can be assessed via the *parameter utilization* metric

$$\mu = \frac{\dim \text{Hom}^{\text{CPTP}}(R^{\text{in}}, R^{\text{out}})}{\dim \text{Hom}_G^{\text{CPTP}}(R^{\text{in}}, R^{\text{out}})}, \quad (17)$$


where $\text{Hom}^{\text{CPTP}}(R^{\text{in}}, R^{\text{out}})$ denotes the set of CPTP maps between \mathcal{B}^{in} and \mathcal{B}^{out} and $\text{Hom}_G^{\text{CPTP}}(R^{\text{in}}, R^{\text{out}})$ its $(G, R^{\text{in}}, R^{\text{out}})$ -equivariant subspace. Note that, $\dim \text{Hom}^{\text{CPTP}}(R^{\text{in}}, R^{\text{out}}) = |\mathcal{H}^{\text{in}}|^2 |\mathcal{H}^{\text{out}}|^2 - |\mathcal{H}^{\text{in}}|^2$ [116]. That is, the larger μ is, the larger is the benefit of using an EQNN, in the sense that available parameters are used more effectively. For instance, by imposing $\mathbb{SU}(2)$ -equivariance on 2-to-2-qubit channels, one reduces the number of free parameters to less than or equal to 14 (see Sec. VI), yielding a reduction of $\mu \geq 240/14 \approx 17$.

C. Intermediate representations as hyperparameters

Here, let us discuss an aspect of EQNNs that has been purposely overlooked up to this point. Namely, while the input representation R^{in} is fixed by the action of the symmetry group on the input data, the intermediate and output representations acting on the spaces \mathcal{B}^l are not. This means that there exists freedom in choosing a sequence of representations $(R^{\text{in}}, R^1, \dots, R^{\text{out}})$ under which the layers are equivariant. That is, we have the following.

Definition 8 (Layered EQNN). An L -layered G -equivariant QNN is defined by a sequence of $L+1$ representations of G , $(R^{\text{in}}, R^1, \dots, R^{\text{out}})$, and a sequence of (G, R^l, R^{l+1}) -equivariant layers.

The equivariance encoded in the L -layered EQNN of the definition given in Eq. (8) can be visualized via the following commutative diagram:



The foregoing discussion allows us to see that $\mathcal{N} = \mathcal{N}_{\theta_L}^L \circ \dots \circ \mathcal{N}_{\theta_1}^1$ is an $(G, R^{\text{in}}, R^{\text{out}})$ -equivariant QNN. Evidently, if we follow such a QNN with (G, R^{out}) -equivariant measurements, we achieve a G -invariant model.

Note that, as previously discussed, a representation defines a Fourier space, meaning that it determines the space over which a layer of EQNN can act or, alternatively, the information in the states that can be accessed. As such, one can use intermediate representations to change how the model accesses and processes information, which can fundamentally determine the success of the learning model.

The most general way of fully specifying a representation is via the multiplicities of the irreps. Thus, the irrep multiplicities m_λ^l of the intermediate representations can be understood as hyperparameters of the EQNN, similar to the number of channels in a conventional CNN. While, in general, there are no strict rules on what representations to use, here we discuss strategies to choose the intermediate representations that are physically meaningful and ease calculations of equivariant layers.

First, one should choose representations that are *natural* on quantum systems. For example, on the space of n qubits $\mathcal{H} = (\mathbb{C}^2)^{\otimes n}$, the unitary group $\mathbb{U}(2)$ has a natural representation that consists of the identical action on each local subsystem via $R(U) = U^{\otimes n}$. Similarly, the cyclic group \mathbb{Z}_n has a natural action on n qubits corresponding to the cyclic shifting of the qubits, captured by the representation $R(g^t) \otimes_{j=1}^n |\psi_j\rangle = \otimes_{j=1}^n |\psi_{j+t \bmod n}\rangle$. Second, the following proposition asserts that equivalent intermediate representations yield the same model expressibility [31,118] and hence it suffices to consider inequivalent ones when designing EQNNs. We defer the proof to Appendix C.

Proposition 2 (Insensitivity to equivalent representations). Consider an EQNN as defined in Definition 8. Then, changing an intermediate representation, R^l , to another representation equivalent to it, VR^lV^\dagger , where V is a unitary, does not change the expressibility of the EQNN.

Finally, we note that the case of finite groups and regular representations (i.e., when the intermediate representations are chosen to be $R_{\text{reg}} : G \rightarrow \mathbb{C}[G]$ corresponding to the

group action on its own group algebra) has been studied in the classical literature under the name of *homogeneous ENNs* [14]. In this case, any equivariant map is a group convolution [11], which can be realized as a unitary operator embedding the classical convolution kernel by the quantum algorithms in Ref. [119]. Combining this with quantum algorithms for polynomial transformations of quantum states [120,121] allows one to quantize classical homogeneous ENNs. In other words, classical homogeneous ENNs can be implemented on a quantum computer as a special case of EQNNs.

D. Field guide to equivariant layers

As previously discussed, intermediate representations can be considered as hyperparameters for the EQNN. In what follows, we define and characterize different types of equivariant layers arising from different intermediate representations.

1. Standard, embedding, and pooling

We start by presenting a definition that categorizes equivariant layers based on the sizes of input and output representations.

Definition 9 (Equivariant layers: standard, embedding, and pooling). Let $\mathcal{N}_{\theta_l}^l : \mathcal{B}^{l-1} \rightarrow \mathcal{B}^l$ be an (G, R^{l-1}, R^l) -equivariant layer. We say that $\mathcal{N}_{\theta_l}^l$ is a pooling layer if $\dim(\mathcal{B}^l) < \dim(\mathcal{B}^{l-1})$, an embedding layer if $\dim(\mathcal{B}^l) > \dim(\mathcal{B}^{l-1})$, and a standard layer if $\dim(\mathcal{B}^l) = \dim(\mathcal{B}^{l-1})$.

Definition 9 does not require the layer to be a quantum channel and is thus applicable beyond the context of quantum to quantum layers, e.g., in quantum algorithms with classical postprocessing, as we discuss in Appendix A. For the special case of EQNNs mapping from a Hilbert space of n qubits to a Hilbert space of m qubits, we say that $\mathcal{N}_{\theta_l}^l$ is a pooling layer if $m < n$, an embedding layer if $m > n$, and a standard layer if $m = n$.

Equivariant quantum circuits have been proposed and used in previous works [45–48,52] mostly in the context of graph problems. However, we note that these fall into *standard layers* and that our framework provides more flexibility, as the operations need not be unitary. An idea of *pooling layers* has been proposed in Ref. [23], although the pooling layer that the authors have used was not equivariant to the symmetry of the considered classification task (see Appendix A). We will provide examples of pooling equivariant layers in later sections. While *embedding layers* have been used to map classical data to quantum data, they are usually not equivariant [38,113] (with a few notable recent exceptions [45,122]), meaning that all the symmetry properties of the classical data are lost during the encoding to quantum states. In addition, to our knowledge, embedding layers mapping quantum data to quantum data

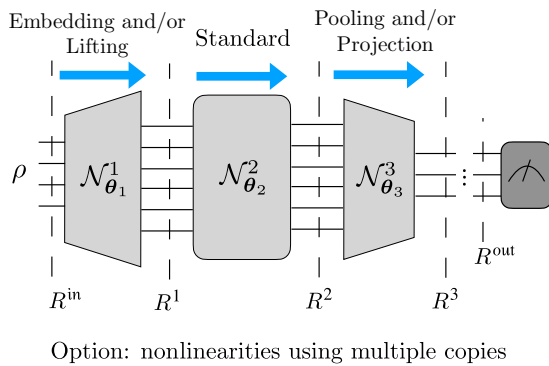


FIG. 3. Different types of equivariant layers in a general architecture of EQNNs. A standard layer maps data between spaces of the same dimension. An embedding (pooling) layer maps the data to a higher-dimensional (smaller-dimensional) space. In a lifting layer, $\ker(R^{l-1}) > \ker(R^l)$, while in a projection layer, $\ker(R^{l-1}) < \ker(R^l)$.

have not been formalized in QML prior to this work. Intuitively, embedding layers equivariantly embed the quantum data into a larger Hilbert space, allowing access to higher-dimensional irreps, and perform nonlinearities (discussed below). A prototypical general EQNN architecture using these equivariant layers inspired by the classical literature [3,14] is illustrated in Fig. 3.

2. Projection and lifting

Another common technique in the classical geometric deep-learning literature is to relax the symmetry constraints in the later layers, typically corresponding to greater-scale features, of the ENNs. This is achieved by *projection layers* (also called reduction layers in some work [14]), which go from representations R^{in} to R^{out} with $\ker(R^{\text{in}}) < \ker(R^{\text{out}})$. Recall that the kernel of a representation R is defined as $\ker(R) := \{g \in G | R(g) = \mathbb{1}\}$, so that R is faithful if and only if $\ker(R) = \{e\}$. Similarly, one can also define *lifting layers*, where $\ker(R^{\text{in}}) > \ker(R^{\text{out}})$. These lifting layers are used as the first layer in many homogeneous ENN architectures [3,8,11,12] but their usefulness is not known in general nonhomogeneous ENNs [14]. Here, we similarly define projection and lifting equivariant layers for EQNNs based on the kernels of the representations as follows.

Definition 10 (Equivariant projection and lifting layers). A (G, R^{l-1}, R^l) -equivariant layer is defined as a projection layer if $\ker(R^{l-1}) < \ker(R^l)$ and a lifting layer if $\ker(R^{l-1}) > \ker(R^l)$.

Projections layers usually become necessary in pooling layers in the presence of outer symmetries that exchange subsystems (see Definition 5), such as \mathbb{Z}_n , D_n , or S_n under qubit-permuting representations. In contrast to *inner symmetries*, which act locally or globally as general

unitaries, such as $\mathbb{U}(2)$ with $R(g) = g^{\otimes n}$, outer symmetries typically have no faithful representations when the number of qubits is reduced by a pooling layer. Hence, in this case it is convenient to use a nonfaithful representation on the output Hilbert space of fewer qubits, i.e., a projection layer. We provide examples of projection layers in Sec. VI.

Lifting layers instead can potentially be beneficial when the symmetry of the problem is unsubstantial and does not greatly reduce the number of free parameters in the model, leading to too expressive EQNNs with potential trainability issues. By lifting to a larger group, one can further reduce the expressibility and potentially improve trainability [31,96]. However, the actual benefit of lifting layers is not known.

Lastly, we note another interpretation of lifting and projection layers as follows. A nonfaithful representation R of G with $\ker(R) = H$ can be thought of as a faithful representation of the quotient group G/H . Then, lifting layers map from a faithful representation of a quotient group to that of a larger quotient group while projection layers have the opposite effect.

3. Nonlinearities

Finally, it is a common practice in QML to assume repeated access to the data set, which means that one can potentially access multiple copies of the input state ρ . The mapping of the form $\rho \rightarrow \rho^{\otimes k}$, which could be applied in the first or an intermediary layer of an EQNN, can thus serve as a *nonlinear equivariant embedding layer*, where $R^{\text{out}} = (R^{\text{in}})^{\otimes k}$.

Definition 11 (Nonlinear equivariant embedding layers). An order- k equivariant nonlinearity in EQNNs is defined as the composition of a map adding $k - 1$ copies of the input state $\phi_{\text{nonlinear}}: \rho \rightarrow \rho^{\otimes k}$.

From the Fourier-space perspective, this operation is analogous to the widely used irrep tensor-product nonlinearity in classical ENNs. For instance, the Clebsch-Gordan decomposition, which computes the tensor product of $\mathbb{SO}(3)$ irreps, has been used in the classical literature to achieve universal nonlinearity [5–7,53]. In the quantum setting, on the other hand, the tensor product is performed *naturally* by composing systems, giving opportunities for equivariant data processing on high-dimensional irreps. Indeed, the first step in the quantum enhanced experiment model [21] performs this nonlinear equivariant layer. Doing so can drastically simplify nonlinear learning tasks [21,44] (see Appendix A).

V. METHODS FOR CONSTRUCTING EQUIVARIANT LAYERS FOR EQNNs

In this section, we describe methods to construct and train layers in EQNNs. Our first step will be to identify,

given a group and its in-and-out representations, the space of equivariant maps. For this purpose, we present three distinct approaches based on finding the null space of a system of matrix equations, the twirling technique, and on the Choi operator. Once the space of equivariant maps is determined, we discuss how to parametrize and optimize over them. An overview of the results in this section can be found in Table II.

A. Simplifying the task of finding equivariant maps

As per Definition 6, a linear map ϕ is equivariant if it satisfies the superoperator equation,

$$\phi \circ \text{Ad}_{R^{\text{in}}(g)} - \text{Ad}_{R^{\text{out}}(g)} \circ \phi = 0, \quad \forall g \in G. \quad (18)$$

The set of all such maps forms a vector space and therefore to characterize them all it suffices to find a basis of this space. While naively it would seem that one needs to solve Eq. (18) for every $g \in G$, we will now see that it is usually enough to solve this equation only over a *well-chosen* subset of elements of the group (or of its Lie algebra for Lie groups of symmetries).

1. Finite groups

We first consider the case in which G is a finite group. Here, we recall the concept of a *generating set*. A subset $S = \{g_1, \dots, g_{|S|}\} \subset G$ is a generating set of G if any element of the group can be written as a product of elements in the generating set. Denoting $\langle S \rangle$ as the closure of S , i.e., the repeated composition of its elements, we say

that S generates G if $\langle S \rangle = G$. For example, the symmetric group S_n can be generated by the set of transpositions. It is a well-known fact in group theory that a finite group can be generated with a subset S of, at most, size $\log_2(|G|)$ [102]. Thus, even exponentially large groups can be handled efficiently through their generating set. In particular, we can simplify the task of finding the equivariant maps via the following theorem.

Theorem 4 (Finite group equivariance). Given a finite group G with generating set S , a linear map ϕ is $(G, R^{\text{in}}, R^{\text{out}})$ -equivariant if and only if

$$\phi \circ \text{Ad}_{R^{\text{in}}(g)} - \text{Ad}_{R^{\text{out}}(g)} \circ \phi = 0, \quad \forall g \in S. \quad (19)$$

2. Lie groups

While Theorem 4 is useful when the group G is finitely generated, many relevant groups, such as the Lie group $\mathbb{U}(d)$, are not. However, we can consider generating sets but now at the Lie-algebra level. In Ref. [60], the authors provide a method for imposing equivariance under Lie groups, where the equivariance constraint is imposed over a basis of the Lie algebra. Evidently, this becomes impractical in the case of large Lie groups, since the method scales linearly on its dimension. Instead, as we prove below, it suffices to impose the constraint only over a generating set. That is, we can consider $s = \{a_1, \dots, a_{|s|}\} \subset \mathfrak{g}$ a generating set for \mathfrak{g} if its Lie closure $\langle s \rangle_{\text{Lie}}$, the repeated nested commutators of the elements of the set, spans the whole Lie algebra. With these concepts at hand, we are ready to impose equivariance at the algebra level.

TABLE II. An overview of the different methods for finding equivariant channels. *Generating set size* denotes that size of the generating set for which the method is better suited. *Main technique* indicates the tools used to create the equivariant channels. Null space uses a linear-algebraic approach to impose equivariance on the generating set of the group or its algebra. Twirling uses the twirl formula defined in Eq. (25) or Eq. (26). The Choi operator block parametrizes the Choi operator via an irrep decomposition. *Time complexity* denotes the computational complexity of the method. The time complexity of Gaussian elimination is $O(d^3)$, where d is the size of the linear system. Assuming that the generating set has size $O(1)$, then the linear system obtained in the null-space method is of size $2^{2(m+n)}$, where n and m are the numbers of qubits at the input and output of the map. For twirling and the Choi operator, the time complexity is dominated by the Haar-integral and irrep decomposition, respectively. In the case of twirling, it can be computed analytically, approximately, or implemented in circuit depending on the problem at hand. *Locality*: determines whether we can control the locality of the operations that we need to implement. *CPTP* indicates whether the output channel is CPTP and how hard it is to impose this condition on the output maps. In null space, it is trivial to impose TP on the solution but imposing CP is more involved. Twirling guarantees CPTP as long as the channel that we twirl is CPTP. In the Choi operator, imposing CP is straightforward but TP might be more involved due to the dimension mismatch introduced by the irrep decomposition. *Kraus rank* indicates whether we can control the Kraus rank of the channel. *Notes* denotes whether we can find a basis for the equivariant map vector space or if we find one map at a time.

Methods	Generating set size	Main technique	Time complexity	Locality controlled?	CPTP condition	Kraus-rank controlled?	Notes	Examples (Sec. VI)
Null space	Small	Linear-algebra based	Gaussian elimination $O(2^{6(m+n)})$	Yes	TP easy	Yes	Find all equivariant maps	SU(2)
Twirling	Any	Weingarten calculus and/or in-circuit twirling	Depends, mainly Haar-integral	Often nontrivial	Yes	Nontrivial	One channel at a time	$\mathbb{Z}_2 \times \mathbb{Z}_2$ Z_n
Choi operator	Any	Irrep decomposition	Depends, mainly irrep decomposition	Nontrivial	CP easy	Yes	Find all equivariant maps	SU(2)

Theorem 5 (Lie group equivariance). Given a compact Lie group G with a Lie algebra \mathfrak{g} generated by s such that exponentiation is surjective, a linear map ϕ is $(G, R^{\text{in}}, R^{\text{out}})$ -equivariant if and only if

$$\text{ad}_{r^{\text{out}}(a)} \circ \phi - \phi \circ \text{ad}_{r^{\text{in}}(a)} = 0, \quad \forall a \in s, \quad (20)$$

where r^{in} and r^{out} are the representations of G induced by R^{in} and R^{out} .

We note that the assumption of surjectivity of the exponential map can be relaxed with the incorporation of additional constraints. This relaxation, together with pertinent examples and proofs of the theorems, are given in Appendix C.

B. Null space, twirling, and Choi operator

With the foregoing discussions in mind, in this section we present three different techniques that can be used to determine equivariant channels.

1. Null-space method

In the *null-space* method, we formulate the equivariance constraints as a linear system of matrix equations, one per element in the generating set, and then solve for their joint null space. This yields a basis for the vector space of equivariant linear maps (not necessarily quantum channels). For the rest of this section, we assume we have a finite group and a set of generators at the group level. The case of Lie groups follows analogously by working at the level of the Lie algebra.

Our method generalizes those in Refs. [60,123] and proceeds as follows. The first step is to represent the superoperators in Eq. (19) as matrices, sometimes referred to as transfer matrices. This can be achieved through the following map $\phi \mapsto \bar{\phi} = \sum_{i,j} \phi_{i,j} |P_i\rangle\rangle\langle\langle P_j|$, where P_j and P_i are Pauli operators in the input and output Hilbert spaces, respectively [124]. Here, $\bar{\phi}$ is a $\dim(\mathcal{B}^{\text{out}}) \times \dim(\mathcal{B}^{\text{in}})$ matrix. The latter transforms Eq. (19) into a matrix multiplication equation of the form

$$\bar{\phi} \cdot \overline{\text{Ad}}_{R^{\text{in}}(g)} - \overline{\text{Ad}}_{R^{\text{out}}(g)} \cdot \bar{\phi} = 0, \quad \forall g \in S. \quad (21)$$

Next, we will perform a vectorization [125], which maps a matrix into a column vector and allows us to write Eq. (21) as

$$M_g \cdot \text{vec}(\bar{\phi}) = 0. \quad (22)$$

Here, $\text{vec}(\bar{\phi})$ is a $\dim(\mathcal{B}^{\text{in}}) \dim(\mathcal{B}^{\text{out}})$ -dimensional column vector and

$$M_g = (\overline{\text{Ad}}_{R^{\text{in}}(g)})^\top \otimes \mathbb{1}_{\dim(\mathcal{B}^{\text{out}})} - \mathbb{1}_{\dim(\mathcal{B}^{\text{in}})} \otimes \overline{\text{Ad}}_{R^{\text{out}}(g)}, \quad (23)$$

is a $\dim(\mathcal{B}^{\text{in}}) \dim(\mathcal{B}^{\text{out}}) \times \dim(\mathcal{B}^{\text{in}}) \dim(\mathcal{B}^{\text{out}})$ matrix. With the foregoing, we can obtain equivariant maps by computing the intersection of the null spaces of each M_g , i.e.,

$$\text{vec}(\bar{\phi}) \in \bigcap_{g \in S} \text{Null}(M_g). \quad (24)$$

In Fig. 4, we present an example of the null-space method.

Here, let us make several important remarks about the null-space method. First, it is clear that this technique can rapidly become computationally expensive. For example, finding equivariant channels mapping from n qubits to m qubits by solving the null spaces through Gaussian elimination [126] has a complexity of $\mathcal{O}(2^{6(m+n)})$. Second, let us note that the solutions of Eq. (24) will lead to a basis for *all* equivariant linear maps and therefore additional steps would be required to find the subset of physically realizable operations (see Sec. VC). For instance, we can obtain trace-preserving (TP) maps by noting that ϕ is TP if and only if $\bar{\phi}$ contains the term $\dim(\mathcal{H}^{\text{in}})/\dim(\mathcal{H}^{\text{out}}) |\mathbb{1}_{\dim(\mathcal{H}^{\text{out}})}\rangle\rangle\langle\langle \mathbb{1}_{\dim(\mathcal{H}^{\text{in}})}|$ and no other terms mapping to $|\mathbb{1}_{\dim(\mathcal{H}^{\text{out}})}\rangle\rangle$.

In practice, we can significantly reduce the computational complexity of this method by restricting the set of Pauli operators that we need to consider in the input and output spaces. This is particularly useful for inner

(a) $\mathcal{H}^{\text{in}} = \mathbb{C}^2, R^{\text{in}}(e) = \mathbb{1}, R^{\text{in}}(\sigma) = X$ $\mathcal{H}^{\text{out}} = \mathbb{C}^2, R^{\text{out}}(e) = \mathbb{1}, R^{\text{out}}(\sigma) = Z$

$\overline{\text{Ad}}_{R^{\text{in}}(e)} = \begin{pmatrix} 1 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 1 \end{pmatrix}$	$\overline{\text{Ad}}_{R^{\text{out}}(e)} = \begin{pmatrix} 1 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 1 \end{pmatrix}$
$\overline{\text{Ad}}_{R^{\text{in}}(\sigma)} = \begin{pmatrix} 1 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & -1 & 0 \\ 0 & 0 & 0 & -1 \end{pmatrix}$	$\overline{\text{Ad}}_{R^{\text{out}}(\sigma)} = \begin{pmatrix} 1 & 0 & 0 & 0 \\ 0 & -1 & 0 & 0 \\ 0 & 0 & -1 & 0 \\ 0 & 0 & 0 & 1 \end{pmatrix}$

(b)

$$\bar{\phi} \in \text{span} \left\{ \begin{pmatrix} 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 \end{pmatrix}, \begin{pmatrix} 0 & 0 & 0 & 0 \\ 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 \end{pmatrix}, \begin{pmatrix} 0 & 1 & 0 & 0 \\ 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 \end{pmatrix}, \begin{pmatrix} 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 1 \\ 0 & 0 & 0 & 0 \end{pmatrix} \right.$$

$$\left. \begin{pmatrix} 1 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 \end{pmatrix}, \begin{pmatrix} 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 \\ 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 0 \end{pmatrix}, \begin{pmatrix} 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 \\ 1 & 0 & 0 & 0 \end{pmatrix}, \begin{pmatrix} 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 1 \\ 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 \end{pmatrix} \right\}$$

$\phi(\rho) = \text{Tr}[\rho] \frac{\mathbb{1}}{2}$

$\phi(\rho) = \frac{X\rho X + Z\rho Z}{2}$

FIG. 4. An example of the null-space method. We demonstrate how to use the null-space method to determine the space of 1-to-1-qubit $(G, R^{\text{in}}, R^{\text{out}})$ -equivariant quantum channels, with $G = \mathbb{Z}_2 = \{e, \sigma\}$, $R^{\text{in}} = \{\mathbb{1}, X\}$ and $R^{\text{out}} = \{\mathbb{1}, Z\}$. (a) The matrix representation of both in and out *adjoint* representations of the symmetry group. (b) A basis for the eight-dimensional solution space, as well as two possible equivariant channels: $\phi(\rho) = \text{Tr}[\rho]/2$ obtained from the solution in red and $\phi(\rho) = (X\rho X + Z\rho Z)/2$ obtained by combining the two solutions in green.

symmetries, where the action of the group can be locally studied. For example, consider the following lemma.

Lemma 2 (Global equivariance via local equivariance). Let $\mathcal{B}^{\text{in(out)}}$ be composite input (output) spaces of the form $\mathcal{B}^{\text{in(out)}} = \bigotimes_j \mathcal{B}_j^{\text{in(out)}}$. Then, assume that the representations acting on each of these spaces takes a tensor-product structure over subsystems as $R^{\text{in(out)}}(g) = \bigotimes_j R_j^{\text{in(out)}}(g)$. For local equivariant channels mapping between each pair of in-and-out subsystems $\phi_j : \mathcal{B}_j^{\text{in}} \rightarrow \mathcal{B}_j^{\text{out}}$ that are $(G, R_j^{\text{in}}, R_j^{\text{out}})$ -equivariant, we have that $\bigotimes_j \phi_j$ is $(G, R^{\text{in}}, R^{\text{out}})$ -equivariant.

Thus, we can find equivariant maps locally and take tensor products of them to obtain a global equivariant layer. While such an approach can greatly reduce the computational cost (e.g., solving for 2-to-1-qubit maps requires dealing with 64×64 matrices), this will come at the cost of expressibility, as the composition of local equivariant channels may have a restricted action when compared to a general equivariant global channel [88].

In the case of outer symmetries such as $G = S_n$ and $R^{\text{in}}(g) = R^{\text{out}}(g) = R_{\text{qub}}(g)$ (as defined in Table I), we can use a generating set S including only local transpositions (i.e., involving only two-body operators). Thus, if we want to obtain maps ϕ containing only one- and two-body terms, we only need to consider the sub-block of M_g corresponding to one- and two-body Pauli operators, reducing its size from exponential to only polynomial.

2. Twirling method

We now explain a second method for finding equivariant maps, based on *twirling*. This approach was first proposed in Ref. [45] to determine equivariant unitary channels. Here, we extend this framework to general nonunitary quantum channels with (possibly) different representations in the input and output spaces of ϕ .

Starting with a given channel $\phi : \mathcal{B}^{\text{in}} \rightarrow \mathcal{B}^{\text{out}}$, we define its twirl over a finite symmetry group G as

$$\mathcal{T}_G[\phi] = \frac{1}{|G|} \sum_{g \in G} \text{Ad}_{R^{\text{out}}(g)} \circ \phi \circ \text{Ad}_{R^{\text{in}}(g)}^\dagger. \quad (25)$$

For the case of Lie groups, we replace the summation with an integral over the Haar measure

$$\mathcal{T}_G[\phi] = \int_G d\mu(g) \text{Ad}_{R^{\text{out}}(g)} \circ \phi \circ \text{Ad}_{R^{\text{in}}(g)}^\dagger. \quad (26)$$

From the invariance of the Haar measure $d\mu$, it is clear that for all ϕ , $\mathcal{T}_G[\phi]$ is $(G, R^{\text{in}}, R^{\text{out}})$ -equivariant. Furthermore, $\mathcal{T}_G[\phi] = \phi$ for all equivariant ϕ . Combining these observations, one can see that \mathcal{T}_G is the projection onto the space

of equivariant maps. This realization allows us to write any channel ϕ as

$$\phi = \mathcal{T}_G[\phi] + \phi_A, \quad (27)$$

where ϕ_A is the ‘‘antisymmetric’’ part of ϕ , i.e., the part satisfying $\mathcal{T}_G[\phi_A] = 0$. As such, any measure of the form $\|\phi_A\|$ quantifies how symmetric ϕ is.

On the practical side, twirling is easy for small groups, as one can efficiently evaluate Eq. (25) (for an example, see Fig. 5). However, for large finite groups or for Lie groups, a direct computation of the twirling becomes cumbersome, requiring the use of more advanced techniques. In Appendix D, we discuss different approaches to implement Eqs. (25) and (26). These range from analytical methods based on the *Weignarten calculus* [127,128] (which requires knowledge of the commutant of the representations) to experimental schemes such as *in-circuit* twirling and *approximate* twirling approaches [129]. In particular, we present two approaches for in-circuit twirling based on either the use of ancilla qubits or classical randomness. Both of these are exemplified in Fig. 5.

For completeness, let us compare the twirling method to the null-space approach. First, we note that one of the main advantages of twirling is that, unlike in the null-space method, we are guaranteed that the twirl of a CPTP channel is also CPTP. However, while the null-space method allows us to find all equivariant maps, twirling is performed one map at a time, meaning that finding a complete basis for the equivariant map vector space could be more intricate (although still possible, as we will show in Sec. VI). As such, if one wants a single equivariant channel, twirling is strongly recommended.

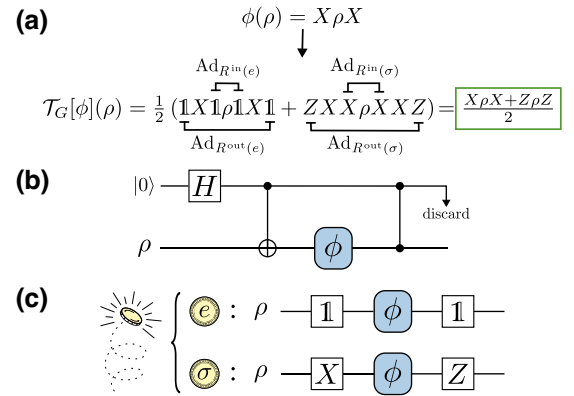


FIG. 5. An example of the twirling method. We demonstrate how to use the twirling method to determine the space of 1-to-1-qubit $(G, R^{\text{in}}, R^{\text{out}})$ -equivariant quantum channels, with $G = \mathbb{Z}_2 = \{e, \sigma\}$, $R^{\text{in}} = \{\mathbb{1}, X\}$ and $R^{\text{out}} = \{\mathbb{1}, Z\}$. (a) An explicit calculation using the twirling formula of Eq. (25). (b) The ancilla-based scheme for in-circuit twirling. (c) The classical-randomness scheme for in-circuit twirling. Both of the schemes in (b) and (c), detailed in Appendix D, recover the twirling in (a).

3. Choi-operator method

Here, we present a third method for finding equivariant channels. We recall from Eq. (15) that the Choi operator of a channel ϕ is given by $J^\phi = \sum_{i,j} |i\rangle\langle j| \otimes \phi(|i\rangle\langle j|)$. Then, as indicated by Lemma 1, the channel will be equivariant if $J^\phi \in \text{comm}(R^{\text{in}*} \otimes R^{\text{out}})$ (where the asterisk (*) denotes the complex conjugate and R^* is the so-called *dual* representation of R) or, alternatively, if for all $g \in G$,

$$J^\phi (R^{\text{in}}(g)^* \otimes R^{\text{out}}(g)) - (R^{\text{in}}(g)^* \otimes R^{\text{out}}(g))J^\phi = 0. \quad (28)$$

While we can vectorize Eq. (28) and obtain equivariant maps by solving for the null space of the matrices $(R^{\text{in}}(g)^* \otimes R^{\text{out}}(g))^\top \otimes \mathbb{1}_{\dim(\mathcal{B}^{\text{out}})} - \mathbb{1}_{\dim(\mathcal{B}^{\text{in}})} \otimes (R^{\text{in}}(g)^* \otimes R^{\text{out}}(g))$, this would not be significantly different from the null-space method previously presented.

Instead, here we focus on a different technique based on the fact that, since $R^{\text{in}*} \otimes R^{\text{out}} = R$ is a valid representation of G , then it has some irrep decomposition $R(g) \cong \bigoplus_q R_q(g) \otimes \mathbb{1}_{m_q}$ (see Theorem 3) and the Choi operator of any equivariant map takes the form

$$J^\phi \cong \bigoplus_q \mathbb{1}_{d_q} \otimes J_q^\phi. \quad (29)$$

Equation (29) allows us to build equivariant maps by controlling precisely how the associated Choi operator acts on each irrep component of the quantum states. We exemplify this method In Fig. 6.

Just as in the null-space method, this approach produces general equivariant linear maps and, hence, additional constraints need to be imposed to find the subset of physical channels. For instance, we can impose TP by requiring that $\text{Tr}_{\text{in}}[J^\phi] = \mathbb{1}_{\dim(\mathcal{H}^{\text{out}})}$, where Tr_{in} indicates the partial trace over \mathcal{H}^{in} . Then, we know that ϕ will be completely positive (CP) if and only if $J^\phi \geq 0$. The last condition implies

FIG. 6. An example of the Choi-operator method. We demonstrate how to use the Choi-operator method to determine the space of 1-to-1-qubit $(G, R^{\text{in}}, R^{\text{out}})$ -equivariant quantum channels, with $G = \mathbb{Z}_2 = \{e, \sigma\}$, $R^{\text{in}} = \{\mathbb{1}, X\}$ and $R^{\text{out}} = \{\mathbb{1}, Z\}$. (a) Isotypic decomposition of the group representation. (b) We show that a specific choice for the block-diagonal components of J^ϕ leads to the map $\phi(\rho) = (X\rho X + Z\rho Z)/2$.

that J^ϕ can be further expressed as [83]

$$J^\phi \cong \bigoplus_q \mathbb{1}_{d_q} \otimes w_q^\dagger w_q, \quad (30)$$

where $w_k \in \mathbb{C}^{m_q \times m_q}$. Moreover, the TP conditions lead to $\sum_q \text{Tr}_{\text{in}}[\mathbb{1}_{d_q} \otimes w_q^\dagger w_q] = \mathbb{1}_{\dim(\mathcal{H}^{\text{out}})}$. Thus, given the irrep decomposition of R , one can construct a basis of CP maps in the block-diagonal form of Eq. (30) and impose the trace-preserving condition afterward [note, however, that taking the partial trace is now more involved due to the subspace mismatch introduced by the isomorphism in Eq. (29)]. Finally, we remark that one could even go a step further and consider conditions for ϕ to be an *extremal* equivariant CPTP channel [130]. Conditions for such, however, are much more involved [83].

Here, let us remark that the main limitation of the Choi-operator approach is that identifying the isomorphism in the irrep decomposition of $R^{\text{in}*} \otimes R^{\text{out}}$ can be in general challenging. For common compact Lie groups, these decompositions are conveniently implemented in a variety of software packages [131,132]. That being said, this method is best suited for local channels, since the size of the Choi operator scales as $\dim(\mathcal{H}^{\text{out}}) \dim(\mathcal{H}^{\text{in}})$. Thus, if $\dim(\mathcal{H}^{\text{out}}) \dim(\mathcal{H}^{\text{in}})$ is not prohibitively large, one can solve for the change of basis of the isotypic decomposition and identify maps with specific irrep actions.

Lastly, implementing the null-space method requires representing channels as matrices. (This could be done for twirling as well.) To check that these maps are actually channels, i.e., CPTP, we may want to convert from a matrix to the Choi operator, for which CPTP is readily verified. We discuss how to perform this conversion in Appendix E 2.

C. Parametrizing the layers of an EQNN

In GQML, we are not only interested in finding equivariant maps but we also want to parametrize and optimize over them. In this section, we show how one can parametrize the layers of an EQNN. For simplicity, we first consider the case of unitary channels and then study the case of general maps. An overview of the methods proposed in this section can be found in Fig. 7.

1. Parametrizing equivariant unitaries

Here, let us consider the case of a unitary EQNN layer with the same input and output representations. That is, $\mathcal{H}^{\text{in}} = \mathcal{H}^{\text{out}}$, $R^{\text{in}} = R^{\text{out}} = R$, and $\mathcal{N}_{\theta_l}^l(\rho) = U_l(\theta_l)\rho U_l(\theta_l)^\dagger$. While this case has been considered in Refs. [44,45,97], we will here review it for completeness.

The simplest way to parametrize a unitary is by expressing it as the exponential of some Hermitian operator usually known as a *generator*, i.e., $U_l(\theta_l) = e^{-i\theta_l H_l}$, where θ_l

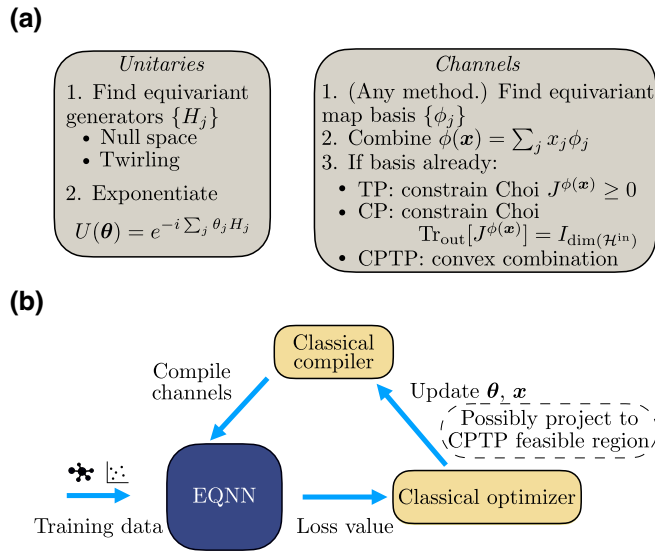


FIG. 7. The procedure to parametrize and optimize equivariant quantum neural networks. (a) We provide techniques to parametrize both equivariant unitaries and general equivariant channels. (b) Once we have a parametrized EQNN, we can proceed to train it. Here, we feed the training data into the EQNN, the outputs of which are used to compute the loss function. Leveraging a classical optimizer, we find updates for the parameters in the EQNN. In the case of channels, it might be necessary to project the updated map to the feasible CPTP region. Note that we can use classical compilers to transform a linear combination of channels into an sequence of gates that we can implement on a quantum device (Appendix E). The procedure is repeated until convergence is achieved.

is a trainable parameter. Evidently, we can obtain (G, R) -equivariant unitaries by taking equivariant generators, i.e., $H_l \in \text{comm}(R)$. Note that we can find equivariant generators via the null-space or twirling approaches previously detailed. While these methods have been presented for superoperators, they can be straightforwardly adapted to the case of operators. Alternatively, one could use the Choi-operator approach and require the solution to be rank-1 (recall that the rank of the Choi operator is the Krauss rank of the associated channel, with unitaries being Kraus-rank-1 channels).

2. Parametrizing equivariant channels

Here, we describe how to parametrize and optimize over equivariant channels. We assume that a basis of equivariant maps (or at least a subset of this basis) has been found via the null-space or Choi-operator method. As mentioned before, while it is relatively easy to find equivariant maps, these need not be physical channels, as they may not be TP, CP, or either. However, one can still parametrize a set of non-CPTP equivariant maps and optimize over them by appropriately constraining the parameters such that the final map is CPTP.

For instance, when using the Choi-operator method, we know that the CP condition is satisfied if $J^\phi \geq 0$. Hence, one could start with some basis of trace-preserving and trace-annihilating equivariant maps $\{J^{\phi_j}\}$, linearly combine them as $J(\mathbf{x}) = \sum_j x_j J^{\phi_j}$, and optimize the set of real parameters $\mathbf{x} = \{x_j\}$ under the constraint that the eigenvalues of $J(\mathbf{x})$ are non-negative. The latter will yield a region of feasible equivariant quantum channels (for an example, see Sec. VI). Note that during the optimization of \mathbf{x} , the update rule might take us outside of the equivariant region, in which case one needs to project back to the feasible space. We further discuss how such projection can be performed in Appendix E. Finally, we note that while it might not be directly obvious how to implement the ensuing channel, one can always transform it into an implementable sequence of gates, acting on a potentially larger space, via compilation techniques [133–136]. Here, we also remark that in many cases, particularly when the maps act on large-dimensional spaces, finding the eigenvalues of $J(\mathbf{x})$ might be quite difficult. For these scenarios, one can simply optimize over a subset of equivariant channels (i.e., maps that are already CPTP), which can be found via twirling. Here, we are guaranteed that any convex combination of equivariant channels will lie in the feasible region, since CPTP channels form a convex set [83].

An alternative approach to constructing equivariant channels is via the Stinespring dilation picture [116]. In this case, we use the fact that any channel can be written as a unitary operation on a larger space, i.e.,

$$\phi(\rho) = \text{Tr}_E[U(\rho \otimes |e\rangle\langle e|)U^\dagger], \quad (31)$$

where $|e\rangle \in \mathcal{H}^E$ is a fixed reference state on an environment Hilbert space \mathcal{H}^E and where Tr_E denotes the trace over \mathcal{H}^E . If $U(R^{\text{in}}(g) \otimes \mathbb{1}_{\dim(\mathcal{H}^E)}) = (R^{\text{out}}(g) \otimes R^{(E)}(g))U$, $\forall g \in G$, then ϕ is a $(G, R^{\text{in}}, R^{\text{out}})$ -equivariant channel. Here, we can use any of the tools previously discussed to find and parametrize U . This approach has the advantage that by fixing the dimension of the environment, we can look for channels of small Kraus rank, which are easier to implement in practice.

VI. APPLICATIONS

In this section, we exemplify the applicability of the methods presented in the previous sections to design EQNNs.

A. $\text{SU}(2)$ -equivariant QCNN

As a first practical application of the present framework, we propose to generalize standard QCNNs [23] to group-equivariant QCNNs. We start by recalling that QCNN have been successfully implemented for error correction, quantum phase detection [23,137], image recognition [138],

and entanglement detection [24]. QCNNS exhibit several key features that make them promising architecture for the near term, such as having a shallow depth or not exhibiting barren plateaus [139]. Despite these advantages, standard QCNNS need not respect the symmetries of a given task. In what follows, we will show how one can design equivariant layers for QCNNS, thus promoting them to group-equivariant QCNNS.

We consider problems where the symmetry group is $\mathbb{S}\mathbb{U}(2)$. This symmetry appears in tasks where the data arise from certain spin-chain models [140–142] and in tasks related to entanglement measures [42,143,144]. Moreover, we consider the case in which the data correspond to n -qubit quantum states and where the input representation of $G = \mathbb{S}\mathbb{U}(2)$ is $R^{\text{in}}(g) = R_{\text{tens}}(g) = g^{\otimes n}$. For ease of implementation on quantum hardware, we restrict ourselves to channels with locality constraints (see Lemma 2). That is, as illustrated in Fig. 8, we want to build an equivariant QCNN that is composed of alternating layers of 2-to-2 standard equivariant channels acting on neighboring qubits and 2-to-1 equivariant pooling channels (for completeness, the 1-to-2-qubit equivariant embedding maps are presented in Appendix F). Of course, this choice of architecture trades locality at the cost of expressibility, as there may be more general equivariant channels on n qubits. However, the success of models with locality constraints [23,139] suggests that this may be an interesting regime regardless.




FIG. 8. $\mathbb{S}\mathbb{U}(2)$ -equivariant QCNN. (a) We consider the problem of building 2-to-2 standard equivariant channels and 2-to-1 equivariant pooling channels. In the figure, we present the respective input and output Hilbert spaces, as well as the input and output representation. (b) In an $\mathbb{S}\mathbb{U}(2)$ -equivariant QCNN, we alternate between 2-to-2 channels acting on neighboring qubits and 2-to-1 equivariant pooling channels that reduce the feature space dimension.

1. 2-to-2 layers via Choi operator

Let us commence by studying 2-to-2-qubit maps via the Choi-operator method. Since the input and output representations are $R(g) = g^{\otimes 2}$, the Choi operator must commute with the representation $(g^*)^{\otimes 2} \otimes g^{\otimes 2}$ [see Eq. (28)]. As $\mathbb{S}\mathbb{U}(2)^*$ shares the same irrep structure as $\mathbb{S}\mathbb{U}(2)$, we can find that the Choi operator for any completely positive $\mathbb{S}\mathbb{U}(2)$ -equivariant map takes the form

$$J^\phi = (\mathbb{1}_5 \otimes A) \oplus (\mathbb{1}_3 \otimes B) \oplus C, \quad (32)$$

where A is a non-negative scalar and B and C are complex positive semidefinite matrices of dimensions three and two, respectively. Thus, the space of such CP equivariant maps is $1^2 + 2^2 + 3^2 = 14$ dimensional.

In the special case of 2-to-2 equivariant *unitary* layers, where $\mathcal{N}_{\theta_l}^l : (\mathbb{C}^2)^{\otimes 2} \rightarrow (\mathbb{C}^2)^{\otimes 2}$ and $\mathcal{N}_{\theta_l}^l(\rho) = U_l(\theta_l)\rho U_l(\theta_l)^\dagger$, we know that if $U_l(\theta_l) = e^{-i\theta_l H_l}$, it suffices to use equivariant generators, i.e., such that $[H_l, g^{\otimes 2}] = 0$ for all $g \in \mathbb{S}\mathbb{U}(2)$. Here, we can use the Schur-Weyl duality [44,145], which states that the only possible equivariant operators are $\mathbb{1}$ and SWAP, which correspond to the two elements of the qubit-permutational representation of S_2 . Without loss of generality, we can choose $H_l = \text{SWAP}$ so that $U_l(\theta_l) = e^{-i\theta_l \text{SWAP}}$. Following Lemma 2, we know that if we compose these two-qubit equivariant unitaries as in Fig. 8, the result will be an n -qubit equivariant unitary.

2. 2-to-1 layers via null space

Next, let us focus on finding 2-to-1-qubit channels using the null-space approach. Since $\mathbb{S}\mathbb{U}(2)$ is a Lie group, we will work at the level of the generators of its Lie algebra, $\mathfrak{su}(2) = \text{span}\{X, Y, Z\}$. Given the representations $g^{\otimes 2}$ and g , the associated basis representations of the algebra are $\{\mathbb{1} \otimes X + X \otimes \mathbb{1}, \mathbb{1} \otimes Y + Y \otimes \mathbb{1}, \mathbb{1} \otimes Z + Z \otimes \mathbb{1}\}$ and $\{X, Y, Z\}$. Thus, one needs to simultaneously solve for the null space of the following matrices:

$$\begin{aligned} M_X &= \overline{\text{ad}_{\mathbb{1}X + X\mathbb{1}}}^\top \otimes \mathbb{1}_2 - \mathbb{1}_4 \otimes \overline{\text{ad}_X}, \\ M_Y &= \overline{\text{ad}_{\mathbb{1}Y + Y\mathbb{1}}}^\top \otimes \mathbb{1}_3 - \mathbb{1}_4 \otimes \overline{\text{ad}_Y}, \\ M_Z &= \overline{\text{ad}_{\mathbb{1}Z + Z\mathbb{1}}}^\top \otimes \mathbb{1}_3 - \mathbb{1}_4 \otimes \overline{\text{ad}_Z}. \end{aligned} \quad (33)$$

Solving, we find five superoperators that form a basis for 2-to-1-qubit $(\mathbb{S}\mathbb{U}(2), g^{\otimes 2}, g)$ -equivariant maps. These are

$$\begin{aligned} \phi_1(\rho) &= \text{Tr}[\rho] \frac{\mathbb{1}}{2}, & \phi_2(\rho) &= \text{Tr}[\rho \text{SWAP}] \frac{\mathbb{1}}{2}, \\ \phi_3(\rho) &= \text{Tr}_A[\rho], & \phi_4(\rho) &= \text{Tr}_B[\rho], \\ \phi_5(\rho) &= \sum_{ijk=1}^3 \text{Tr}[\rho \sigma_i \sigma_j] \epsilon_{ijk} \sigma_k. \end{aligned} \quad (34)$$

Notably, in addition to expected $\mathbb{S}\mathbb{U}(2)$ -equivariant maps such as trace, partial traces, and SWAP measurement, we identify a potentially interesting new equivariant map $\phi_5(\rho)$ that is dubbed the *cross-product map* (we further study its properties in Appendix F). Note that ϕ_1 , ϕ_3 , and ϕ_4 are trace preserving while ϕ_5 is trace annihilating. One can also verify that ϕ_2 may nontrivially alter trace. As the only map that may do so, we can drop it from our basis set for being nonphysical. To continue and find the set of equivariant quantum channels, we first make a modification to our basis set. In the Pauli basis, we have $\phi_1 \leftrightarrow 2|\mathbb{1}\rangle\rangle\langle\langle\mathbb{1}, \mathbb{1}|$ and it is easy to see that both ϕ_3 and ϕ_4 also contain this term. Thus, we can remove it, leaving trace-annihilating versions of partial trace, which we will denote by ϕ'_3 and ϕ'_4 . Thus, any TP map must take the form

$$\phi(x, y, z) = \phi_1 + x\phi_5 + y\phi'_3 + z\phi'_4, \quad (35)$$

where the coefficients are real numbers. It remains to find the region such that this channel is CP. This can be done via the Choi operators of these channels. That is, we would like to find

$$\{x, y, z \in \mathbb{R} : J\phi_1 + xJ\phi_5 + yJ\phi'_3 + zJ\phi'_4 \geq 0\}. \quad (36)$$

Note that the coefficients here must be real numbers for the Choi operator of the sum to be positive (as the Choi operators in the sum are linearly independent). Requiring the eigenvalues of this linear combination to be non-negative yields the feasible region

$$\begin{aligned} x, y, z : y + z \leq 1, \quad \text{and} \\ y + z \geq \sqrt{3x^2 + 4(y^2 - yz + z^2)} - 1. \end{aligned} \quad (37)$$

This region is illustrated in Fig. 9. Here, we note that as the set of equivariant channels is convex, this feasible parameter region is a convex subset of \mathbb{R}^3 .

A crucial aspect to note is that when training the $\mathbb{S}\mathbb{U}(2)$ -equivariant QCNN, one can directly train over the coefficients x , y , and z of each pooling channel $\phi(x, y, z)$ of the form in Eq. (35). When training an equivariant QCNN, e.g., using gradient descent, we will obtain parameter updates $(x^{(t+1)}, y^{(t+1)}, z^{(t+1)}) \leftarrow (x^{(t)}, y^{(t)}, z^{(t)}) - \alpha D_t((x^{(t)}, y^{(t)}, z^{(t)}))$. To ensure that the operations remain physical, one would continually solve the projection at each iteration. This can be turned into the convex optimization problem

$$\begin{aligned} \min_{x, y, z} \|(x^{(t+1)}, y^{(t+1)}, z^{(t+1)}) - (x, y, z)\|^2, \\ \text{subject to Eq. (37)} \end{aligned} \quad (38)$$

over a convex domain (see Appendix E).




FIG. 9. The region of parameter space leading to CPTP channels. Using the null-space method, we can find a basis for all 2-to-1 ($\mathbb{S}\mathbb{U}(2), g^{\otimes 2}, g$)-equivariant pooling maps. These can then be linearly combined to form a general parametrized equivariant map as in Eq. (35) and we find in Eq. (37) the region in parameters space leading to CPTP channels. Here, we depict said region as the volume of the hyperbole (red) below the plane (green).

B. Various examples and physical considerations

In this section, we present additional applications of the methods detailed in Sec. V. In particular, these are now applied to discrete groups and motivated from practical problems.

1. $\mathbb{Z}_2 \times \mathbb{Z}_2$ -equivariant layers

Here, we consider problems with $\mathbb{Z}_2 \times \mathbb{Z}_2$ symmetry, which are common in spin chain models such as the $S = 1$ Haldane chain [23,140,146], or in classical data on the two-dimensional plane [45].

We start with a task on n qubits, where the representation of $\mathbb{Z}_2 \times \mathbb{Z}_2$ is given by

$$\begin{aligned} R^{\text{in}}(e, e) &= \mathbb{1}_{2^n}, \quad R^{\text{in}}(\sigma, e) = \prod_{i \in O} X_i \otimes \mathbb{1}_{2^{n/2}}, \\ R^{\text{in}}(e, \sigma) &= \mathbb{1}_{2^{n/2}} \otimes \prod_{i \in E} X_i, \quad R^{\text{in}}(\sigma, \sigma) = \prod_{i \in E \cup O} X_i. \end{aligned}$$

Here, the sets O and E , respectively, contain the odd and even qubit labels. From the foregoing, we are interested in finding all equivariant standard unitary channels where $R^{\text{in}} = R^{\text{out}}$. Since the group is small, we can readily employ the twirling method to determine the set of all equivariant generators. The set of all such generators forms a subalgebra of the unitary algebra $u(2^n)$, which we denote

$u^{\mathbb{Z}_2 \times \mathbb{Z}_2}(2^n)$ and which is given by

$$u^{\mathbb{Z}_2 \times \mathbb{Z}_2}(2^n) = \text{span}_{\mathbb{R}} i \left\{ \begin{array}{l} \text{Pauli strings with even \# of Y's} \\ \text{and Z's on both } E \text{ and } O \end{array} \right\}.$$

Next, let us find pooling channels that are equivariant with respect to $\mathbb{Z}_2 \times \mathbb{Z}_2$. Similarly to the $\text{SU}(2)$ case previously considered, the representations of the group act locally, meaning that we can again consider local 2-to-1 equivariant pooling maps and later combine them into a global equivariant channel (see Lemma 2). On any pair of neighboring qubits, the input representations are $\{\mathbb{1} \otimes \mathbb{1}, \mathbb{1} \otimes X, X \otimes \mathbb{1}, X \otimes X\}$ and we set the output representation acting on a single qubit to be $\{\mathbb{1}, \mathbb{1}, X, X\}$. Note that the output representation is not faithful. As $|\mathbb{Z}_2 \times \mathbb{Z}_2| = 4$, we can again readily apply the twirling procedure. To do so, we begin with a basis for trace-preserving maps; i.e., matrices in the Pauli-string basis of the form $2|\mathbb{1}\rangle\rangle\langle\langle \mathbb{1} \otimes \mathbb{1} + |P_{\text{out}}\rangle\rangle\langle\langle P_{\text{in}}$, where $P_{\text{out}} \neq \mathbb{1}$. A simple counting argument reveals that there are 48 such matrices. By twirling all 48 maps and extracting a linearly independent set, we find the following 13-element basis for 2-to-1 equivariant projective poolings:

$$\varphi_1(\rho) = \text{Tr}[\rho] \frac{\mathbb{1}}{2}, \quad \varphi_2(\rho) = \text{Tr}[\rho] X, \quad (39)$$

$$\varphi_3(\rho) = \text{Tr}[(\mathbb{1} \otimes X)\rho] X, \quad \varphi_4(\rho) = \text{Tr}[(X \otimes \mathbb{1})\rho] X, \quad (40)$$

$$\varphi_5(\rho) = \text{Tr}[(X \otimes X)\rho] X, \quad \varphi_6(\rho) = \text{Tr}[(Y \otimes \mathbb{1})\rho] Y, \quad (41)$$

$$\varphi_7(\rho) = \text{Tr}[(Y \otimes X)\rho] Y, \quad \varphi_8(\rho) = \text{Tr}[(Z \otimes \mathbb{1})\rho] Y, \quad (42)$$

$$\varphi_9(\rho) = \text{Tr}[(Z \otimes X)\rho] Y, \quad \varphi_{10}(\rho) = \text{Tr}[(Z \otimes \mathbb{1})\rho] Z, \quad (43)$$

$$\varphi_{11}(\rho) = \text{Tr}[(Z \otimes X)\rho] Z, \quad \varphi_{12}(\rho) = \text{Tr}[(Y \otimes \mathbb{1})\rho] Z, \quad (44)$$

$$\varphi_{13}(\rho) = \text{Tr}[(Y \otimes X)\rho] Z. \quad (45)$$

2. \mathbb{Z}_n -equivariant layers

We now proceed to analyze a problem with \mathbb{Z}_n symmetry, the representation of which cyclically shifts qubits as $R(g^t) \otimes_{j=1}^n |\psi_j\rangle = \otimes_{j=1}^n |\psi_{j+t \bmod n}\rangle$. Such symmetry arises naturally in condensed matter problems with periodic boundary conditions [25, 147, 148].

Let us start by finding all equivariant standard unitary maps where $R^{\text{in}} = R^{\text{out}}$. Since the group is small, we opt for the twirling approach. For the sake of implementability, we will seek channels composed of one- and two-qubit gates. We can readily see that the twirl of a single-qubit generator such as X_1 leads to the sum of single-qubit

operators $\mathcal{T}[X_1] = 1/n \sum_{j=1}^n X_j$. Similarly, twirling a two-body generator will lead to a sum of two-body generators (e.g., $Z_1 Z_2 \xrightarrow{\mathcal{T}} \sum_j Z_j Z_{j+1}$). Notably, these generators lead to equivariant unitaries of the form $U_l = e^{-i\theta_l(1/n) \sum_{j=1}^n X_j} = \prod_{j=1}^n e^{-i\theta_l X_j/n}$. The foregoing shows a crucial implication of outer symmetries (such as \mathbb{Z}_n): in many cases, equivariance in unitary layers can be achieved by *correlating* parameters of local gates within a layer [97, 149].

Note that a necessary condition for the foregoing to hold is that all the terms in the twirled operator must be mutually commuting. One can see, however, that the twirl of $Z_1 Y_2$ leads to $\sum_j Z_j Y_{j+1}$, which is a global operator the terms of which are noncommuting and that can be challenging to implement on near-term devices. An alternative here is to employ a randomized method. First, we construct unitaries $U_{l,O} = \prod_{j \in O} e^{-i\theta_l Z_j Y_{j+1}}$ and $U_{l,E} = \prod_{j \in E} e^{-i\theta_l Z_j Y_{j+1}}$ (where we recall that O and E , respectively, contain the odd and even qubit labels), which are $\mathbb{Z}_{n/2}$ -equivariant. Then, to achieve \mathbb{Z}_n -equivariance, we can apply either $U_{l,O}$ or $U_{l,E}$ at random and with equal probability, effectively performing the quantum channel $\rho \rightarrow (U_{l,O} \rho U_{l,O}^\dagger + U_{l,E} \rho U_{l,E}^\dagger)/2$. This channel can be readily shown to be \mathbb{Z}_n -equivariant.

This trick of randomly applying local channels can also be applied to equivariant layers that have different numbers of qubits in the input and output. For example, a \mathbb{Z}_n -equivariant projection layer that reduces the number of qubits from n to $n/2$ is $\Phi_{\mathbb{Z}_n} : \rho \rightarrow (\text{Tr}_{\text{odd}} \rho + \text{Tr}_{\text{even}} \rho)/2$. Observe that the output representation has a nontrivial kernel isomorphic to \mathbb{Z}_2 , as the number of qubits is halved. One can readily extend these ideas to projecting pooling layer over other outer symmetry groups G , such as S_n -equivariance: $\Phi_{S_n} : \rho \rightarrow (\sum_S \text{Tr}_S \rho) / \binom{n}{n/2}$, where S are uniformly random subsets of $n/2$ qubits. The Hoeffding's bound implies that only $O(\log |G|)$ samples are needed for this method to converge to within a specified error bound. In a sense, this is similar to the dropout regularization technique in neural networks [150].

VII. NUMERICAL EXPERIMENTS

In this section, we numerically compare the performance of $\text{SU}(2)$ -equivariant QCNNS constructed in Sec. VI A against a problem-agnostic QCNN in a quantum phase-classification task. Similar to the classical problem of assigning the correct labels to images, the task of classifying quantum phases of matter can be carried out in a supervised setting and provides a natural playground to study the efficiency of equivariant quantum learning model.

A. Bond-alternating Heisenberg model

The one-dimensional (1D) XXX Heisenberg model describes the behavior of a 1D chain of spin-1/2 particles

coupled through the standard Heisenberg interaction Hamiltonian between nearest neighbors:

$$H = \sum_i \sum_{k=x,y,z} J_k S_i^k S_{i+1}^k, \quad (46)$$

where S_i is the spin operators for the i th spin, with $S = (S^x, S^y, S^z) = \frac{1}{2}(X, Y, Z)$ and $J_x = J_y = J_z = J$. The bond-alternating XXX Heisenberg model is a generalization of the regular XXX Heisenberg model, in which the exchange coupling constant alternates between two different values J_1 and J_2 :

$$H = J_1 \sum_{i \text{ even}} S_i \cdot S_{i+1} + J_2 \sum_{i \text{ odd}} S_i \cdot S_{i+1}. \quad (47)$$

We consider the model described by Eq. (47) with open boundary conditions and with both the couplings in the ferromagnetic regime, i.e., $J_{1,2} > 0$. In this case, a *trivial* phase and a *topologically protected* phase are defined by the expectation value of the partial reflection many-body topological invariant [141,142,151]. The quantum phase transition between these two phases occurs at a critical value of the bond-alternation parameter $\alpha = J_2/J_1$. When $\alpha < 1$, the system is in the trivial phase; otherwise, the system is in the topologically protected phase.

The Hamiltonian can be readily seen to possess an $\mathbb{S}\mathbb{U}(2)$ symmetry. The symmetry extends to the whole system through the tensor-product representation $R^{\text{in}}(g) = R_{\text{tens}}(g) = g^{\otimes n}$. This is also a symmetry of the phase labels, as quantum phases are global properties of the ground states of the model and symmetries of the Hamiltonian are also symmetries of the ground space. Indeed if $[H, g^{\otimes n}] = 0$ and $|\psi\rangle$ is a ground state, then $g^{\otimes n} |\psi\rangle$ is also a ground state. Thus, a quantum phase classifier can be endowed with this inductive bias using the $\mathbb{S}\mathbb{U}(2)$ -equivariant quantum maps from Sec. VI A. Another inductive bias that we can utilize is translation symmetry: shifting the qubits by two sites leaves the model unchanged. One way to exploit this is parameter sharing within each layer of the $\mathbb{S}\mathbb{U}(2)$ -equivariant QCNN, as discussed in Sec. VI B.

B. Defining the learning models

We now describe the learning model in more details. We use the $\mathbb{S}\mathbb{U}(2)$ -EQCNN architecture in Fig. 8, where each standard (convolution) layer consists of two brickwork (sub)layers of two-qubit $\mathbb{S}\mathbb{U}(2)$ -equivariant gates of the form $U(\theta) = e^{-i\theta \text{SWAP}}$. Due to translation symmetry, we further enable parameter sharing of two-qubit gates within each of such sublayers. This leads to having two parameters for each standard convolution layer. If needed, we can repeat the standard layers more than once before applying the pooling layer.

Finally, we choose an equivariant observable operator \hat{O} to measure at the end of the EQCNN. As discussed in

the previous sections, the final equivariant measurements belong to the commutant of the output representation R^{out} of the last layer of the EQCNN. Now, since we need two outputs to label the two phases $y_{\text{trivial}} = 1$ and $y_{\text{topological}} = 0$, it is tempting to end the EQCNN with a binary measurement on $m = 1$ qubit. However, from the discussion in Sec. VI A, we know that the commutant of the defining representation of $\mathbb{S}\mathbb{U}(2)$ over a single qubit is the trivial set $\text{comm}(R_{\text{natural}}) = \{\mathbb{1}\}$. Thus we choose a EQCNN that ends with $m = 2$ such that the commutant of $g^{\otimes 2}$ contains the nontrivial element SWAP. Conveniently, SWAP has two eigenvalues ± 1 , so that we can bind, say, the $+1$ outcome to y_{trivial} and the -1 one to $y_{\text{topological}}$. We adopt this strategy, with a little modification to have the output of the EQCNN, that we will indicate as $f_{\theta}(\rho)$, take values in $[0, 1]$. Namely, we define

$$f_{\theta}(\rho) = \frac{\text{Tr}[\phi_{\theta}(\rho) \text{SWAP}] + 1}{2}, \quad (48)$$

where ϕ_{θ} denotes the EQCNN that outputs two qubits in the last layer. Then, we assign the predicted phase label to any input state ρ as

$$y_{\theta}(\rho) = \begin{cases} \text{trivial}, & \text{if } f_{\theta}(\rho) > \tau, \\ \text{topological}, & \text{if } f_{\theta}(\rho) < \tau, \end{cases} \quad (49)$$

for some trainable threshold value τ that is initialized to be $\tau = 0.5$.

We test this EQCNN architecture against a QCNN with no inductive biases. In particular, we use a QCNN whose standard layers are inspired by the hardware efficient ansatz (HEA) [152], whereas the pooling layers consist of simple alternate partial traces (at each pooling operation we discard half of the qubits). The classification will then proceed as for the $\mathbb{S}\mathbb{U}(2)$ EQCNN, with a SWAP measurement and phase assignment described in Eq. (49). We dub the nonequivariant QCNN as the HEA QCNN.

C. Training procedure

We use the standard ML pipeline of supervised learning:

- (1) We collect a training data set $\mathcal{D}_{\text{train}}^{N_T}$, where N_T is the size of the data set, by choosing some representative values of the parameter J_2 while always keeping $J_1 = 1$ and then analytically computing the ground states $|\psi\rangle_g^{J_2/J_1}$ of the Hamiltonian in Eq. (47). Knowing the phase diagram of the alternating model, which is shown in Fig. 10, especially that the critical value at which the transition happens $\alpha = J_2/J_1 = 1$, we can then associate these states with their true labels $y \in \{0, 1\}$. In particular, we try a training data set made of $N_T = (2, 4, 6, 8, 10, 12)$ ground states, always distributed homogeneously in




FIG. 10. The 1D bond-alternating XXX Heisenberg model [Eq. (47)] and its phase diagram in terms of exchange coupling constants $J_{1,2} > 0$.

the range $J_2/J_1 \in [0, 1]$. For example, for $N_T = 2$, we use $\mathcal{D}_{\text{train}}^2 = \{(|\psi\rangle_g^{0.25}, 1), (|\psi\rangle_g^{0.75}, 0)\}$.

- (2) We initialize the learning model at hand, equivariant or not, with random parameters θ .
- (3) We select an optimizer for the learning model. In our case, we always use ADAM [153], the gold standard of gradient-based optimization in ML.
- (4) For a number of *epochs* E , we divide the training data set $\mathcal{D}_{\text{train}}^{N_T}$ into batches of size $n_{\text{batch}} = 2$. For each batch, the training states $|\psi_i\rangle$ are processed by the model to output the predicted label $y_\theta(|\psi_i\rangle)$ and the mean-squared-error loss function is computed by comparing the predictions to the real labels y_i :

$$L_\theta = \frac{1}{n_{\text{batch}}} \sum_{i=1}^{n_{\text{batch}}} (y_\theta(|\psi_i\rangle) - y_i)^2. \quad (50)$$

We then compute the gradient of L_θ and use the optimizer to update the parameters of the model. The goal is to minimize L_θ .

- (5) The QCNN outputs for the training states are used to update the threshold τ . In particular, only the two training points that are closer to the critical value $\alpha = 1$ are considered and the threshold value is set to the average of the corresponding outputs.
- (6) As an additional figure of merit for the training, we keep track of the prediction accuracy of the model.
- (7) At the end of the last epoch, we let the model predict the labels of the whole test data set and we compute its final accuracy as a measure of the goodness of the training. Then, we also plot the predicted phase diagram to get a visual proof of the performance of the model.

D. Training results

We are now ready to illustrate the results of our numerics. First, we must state that we have not been able to train the EQCNN when using the general 2-to-1 pooling layers described in Sec. VIA, as the projection step [Eq. (38)] onto the feasible CPTP region seems to cause instability in the optimizing procedure. We leave a full

numerical analysis of equivariant quantum learning models to a future upcoming work and here focus on the more simple tracing pooling operations. That is, the EQCNN architecture is still the one depicted in Fig. 8 but the pooling operations are just 2-to-1 partial trace channels, corresponding to the solution ϕ_3 in Eq. (34). In other words, the $SU(2)$ EQCNN and HEA QCNN use the same pooling layers (but still different convolution layers—equivariant versus HEA). The training results are illustrated in Figs. 11 and 12.

We have considered system sizes ranging from $N = 6$ to $N = 13$ and trained both the EQCNN and the HEA QCNN according to the training loop described in Sec. VII C for a fixed number of training epochs $E = 750$. Since the two architectures are very different and the EQCNN, as opposed to the HEA QCNN, uses parameter sharing, in order to have a fair comparison we have decided to stack multiple standard layers before each pooling one in the EQCNN, in such a way as to have a similar amount of training parameters for both the learning models. In Fig. 11, we show the predicted phase diagrams for $N = 12$ and $N = 13$. The thing that immediately stands out is Fig. 11(d). While the other three plots basically showcase similar behavior, with the QCNN at hand being able to efficiently separate the two phases of the alternating model, Fig. 11(d) shows a cloudy behavior of the HEA-QCNN predictions, as it has assigned different phases even for states with similar parameters α . This is in sharp contrast with the trained EQCNN [Fig. 11(c)], which has successfully learned to classify the two phases with excellent accuracy, demonstrating the advantage of equivariant models.

Interestingly, for the $N = 12$ qubit case [Figs. 11(a) and 11(b)], the EQCNN does not significantly outperform the HEA QCNN. This is due to fact that there is actually no need for equivariance in that case. Indeed, equivariance is meant to enhance the performance of learning models that deal with labels that are invariant under some symmetry group but this invariance should *not* come from the invariance of the input states themselves. Think yet again of the classical problem of classifying images of cats and dogs: the labels, i.e., the semantic meanings of the images, are invariant if we translate the images but the images themselves are not translation invariant. On the other hand, if we translate images that are full of either black or white pixels, instead of showing cats and dogs, the labels (the colors) of the images are translation invariant simply because the images do not change. This is what is happening in our case. We have already discussed that if a Hamiltonian H is symmetric under a group G , any unitary representation of it, U_G , leaves the ground space unchanged. For a nondegenerate ground space, this means that the unique ground state is invariant under the group action, in analogy to the above black and/or white image example, and thus equivariant learning models are not needed. For degenerate

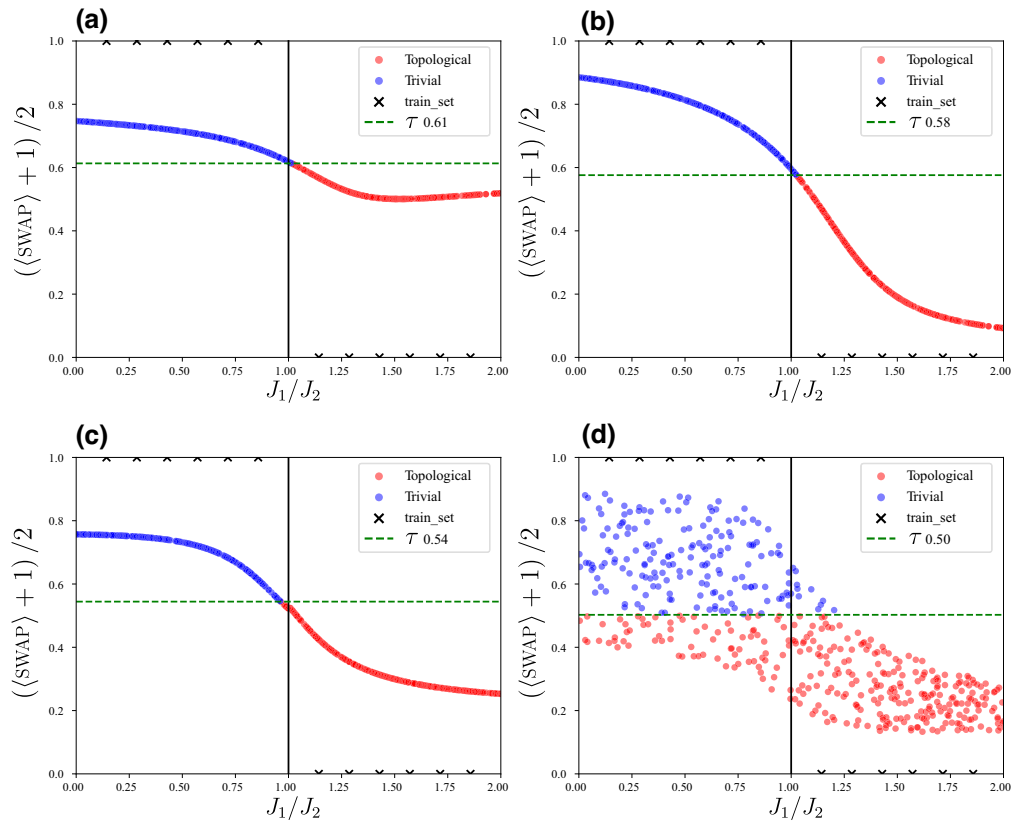


FIG. 11. The predicted phase diagrams. The four panels show the phase diagram of the 1D bond-alternating XXX Heisenberg model for system sizes of $N = 12$ and $N = 13$ qubits, as reconstructed by a trained $\mathbb{S}\text{U}(2)$ EQCNN or HEA QCNN. In particular, each panel shows the QCNN output when it is tested against a data set of 500 homogeneously distributed ground states. States the output of which is above (below) the optimal threshold τ (green dashed line) are colored in blue (red) and classified as belonging to the trivial (topological) phase. The training points are shown as black crosses. The vertical solid black line is the theoretical critical value $J_2/J_1 = 1$. The configurations leading to the panels are the following: (a) $\mathbb{S}\text{U}(2)$ EQCNN, $N = 12$, 60 trainable parameters, 12 training points; (b) HEA QCNN, $N = 12$, 63 trainable parameters, 12 training points; (c) $\mathbb{S}\text{U}(2)$ EQCNN, $N = 13$, 66 trainable parameters, 12 training points; (d) HEA QCNN, $N = 13$, 66 trainable parameters, 12 training points. Details about the training procedure are given in the main text.

ground states, i.e., when the Hamiltonian symmetry is broken, the symmetry group action does change the ground states nontrivially and it rotates them within the ground space. Degenerate ground states are akin to images of cats and dogs and, as such, equivariance can finally shine. Indeed, the alternating model is degenerate for odd system sizes, while for even system sizes the ground state is unique. This explains the different behavior shown in Fig. 11 between $N = 12$ and $N = 13$.

The foregoing discussion also motivates the analysis shown in Fig. 12. There, we show a statistical study of the performance of the EQCNN and the HEA QCNN when tackling even and odd system sizes. As is evident from Fig. 12(a), enforcing equivariance when it is not needed can be more detrimental than beneficial. Indeed, the reduced expressibility of the learning model is not compensated by any benefit and training instabilities emerge, as evidenced by the large error bars in Fig. 12(a). However,

when there is a reason to use equivariance, as there is for the odd-size states studied in Fig. 12(b), the advantage of using the EQCNN against a noninformed one is clear. Already, with only two training points, the equivariant QCNN performs very well, while the HEA QCNN needs more training data to generalize well. Interestingly, even with a minimum number of trainable parameters—that for the system sizes studied ranges from four to six—the EQCNN seems to perform better than the nonequivariant one.

As stated at the outset, this is only a preliminary analysis on a simple learning task and as such we postpone any general conclusions until further studies on the performance of equivariant quantum learning models on more complex systems, against different nonequivariant architectures, and for different symmetry groups. Nonetheless, we think that the preliminary numerical results shown in this section hint at confirming that injecting inductive biases into quantum

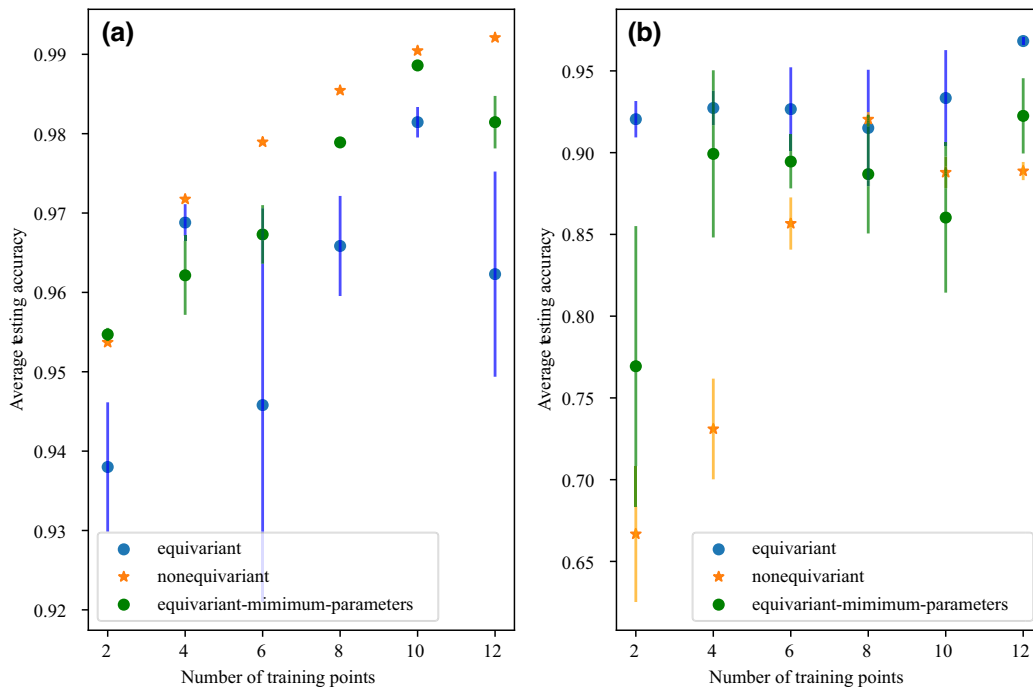


FIG. 12. The actual power of equivariance, showing the mean and variance of the testing accuracy reached by trained QCNNS on the bond-alternating Heisenberg model of (a) even and (b) odd sizes. The average is conducted on both the chosen sizes, (6, 8, 10, 12) for the even case and (7, 9, 11, 13) for the odd one, and on ten different randomly initialized training runs for each problem size. The results are plotted against the number of training data points, N_T . The blue circles refer to the EQCNN with the same number of parameters as the HEA QCNN (orange stars). The green circles describe the EQCNN with the minimum number of parameters possible, i.e., the architecture, as it is, in Fig. 8, with only two standard layers before each pooling layer. These plots demonstrate that equivariance provides significant improvements when, and only when, there is degeneracy in the ground space.

neural networks boosts their performance, paving the way to the design of new more efficiently implementable and trainable variational quantum ML models.

VIII. DISCUSSION AND OUTLOOK

Geometric QML is a new and exciting field that seeks to produce helpful inductive biases for QML models based on the symmetries of the problem at hand. While several proposals already exist in the literature within the field of GQML [44–48,52], these mainly deal with unitary models that maintain the same group representation throughout the computation. In this work, we generalize previous results and we present a theoretical framework to understand, design, and optimize over general equivariant channels, which we refer to as EQNNs. While presented in the setting of supervised learning, our work is readily applicable to other contexts such as unsupervised learning [105,106], generative modeling [107–110], or reinforcement learning [111,112].

Our first main contribution is a characterization of the action of equivariant layers as generalized Fourier actions. We argue that the isotypic decomposition of the representation of the symmetry determines a generalized Fourier space over which the EQNNs act. This realization not only

allows us to characterize the number of free parameters in an EQNN but it also unravels the crucial importance that the choice of representation has. That is, different representations have different block-diagonal structures and, hence, can act on different generalized Fourier spaces and see different parts of the information encoded in the quantum states. Then, we provide a general classification of EQNN layers, introducing the so-called standard, embedding, pooling, projection, and lifting layers, and we note that nonlinearities can be introduced via multiple copies of the data. As a by-product, we highlight the exciting possibility of accessing higher-dimensional irreps of the group symmetry via these nonlinearities, which can be a venue to access information that would otherwise be unavailable.

Our next main contribution is the description of three methods to construct EQNN layers. In the first, which we call the null-space method, we map the equivariance constraints to a linear system of matrix equations and then solve for their joint null space. The second method leverages the technique of twirling over a group, whereby a channel is projected onto the space of equivariant maps. In our third method, we use the Choi operator of the map to create equivariant layers with specific irrep actions. Our methods can find unitary or nonunitary equivariant layers efficiently even when the symmetry group is exponentially

large, rendering applications for groups that are inaccessible using existing methods in prior literature. We then compare the strengths and shortcomings of each method, presenting scenarios in which one should be favored over the other. Our final key contribution is to show how to parametrize and optimize EQNNs. In particular, since our work seeks to find equivariant channels, we discuss how one can guarantee that the ensuing maps are physical and potentially easy to implement. To finish, we exemplify our methods by generalizing standard QCNNS to group-equivariant QCNNS and we show how to create, from the ground up, an $SU(2)$ -equivariant QCNN. Finally, we apply this model to a quantum phase-classification task on the 1D bond-alternating Heisenberg model and numerically demonstrate its superior performance over a symmetry-agnostic QCNN.

A. Equivariance versus barren plateaus, local minima, and data requirements

Here, we argue why EQNNs can alleviate some of the crucial issues in QML, such as barren plateaus, excessive local minima, and poor data requirements.

First, we recall that the barren-plateau phenomenon refers to the exponential concentration of gradients exhibited by certain variational quantum models that result in an exponential flattening of the training landscape and, concomitantly, in an exponential demand of measurement shots to accurately resolve a parameter update [27–41]. The presence or absence of barren plateaus has been directly linked to the expressibility of the model [31,35,36,39], such that highly expressible architectures exhibit smaller gradients. In our context, the imposition of symmetry constraints to the quantum neural network is expected to shrink—in a problem-oriented way—its expressibility, alleviating such gradient-vanishing issues.

Another challenge in training QML models is spurious local minima in the loss landscape. It is known that agnostic models exhibit landscapes that are plagued by local minima [154–157]. However, it is also known that there exists a critical number of trainable parameters above which the model can become overparametrized, meaning that all spurious local traps disappear [158]. While reaching the overparametrization regime requires exponentially deep circuits for agnostic ansatzes, it has been proven that certain architectures (with reduced expressibilities) can be efficiently overparametrized with polynomial depth. Thus, the hope is that by restricting the expressibility of the model via geometric priors, one can reduce and realistically reach the overparametrization threshold, thus getting rid of fake local minima.

Finally, we discuss sample complexity. The ultimate goal of supervised ML is to make predictions on unseen data. This is often characterized by generalization bounds, which measure the difference between the performance of

the learned model on training and testing data. Recent work has studied the training sample complexity needed for QML models to generalize [25,159–161]. In particular, in Ref. [25] it has been shown that the training sample complexity typically scales polynomially with the number of trainable parameters. Given a trainable QNN, we have seen that imposing equivariance can drastically reduce the numbers of free parameters, thus implying that incorporating equivariance allows for stronger (more optimistic) bounds on generalization performance. Further, the bounds of Ref. [25] are statistical and worst-case (over all possible learning tasks), meaning that EQNNs when applied to the corresponding symmetric learning tasks could potentially achieve better generalizations than indicated by these bounds.

While the previous arguments merely indicate *why* equivariance can improve the performance of a model (in terms of trainability and generalization), these do not constitute a proof that equivariance can indeed fulfill these promises. However, we refer the reader to the recent work of Ref. [96] which studies S_n -equivariant models and rigorously proves that the equivariance constraints lead to an architecture that avoids barren plateaus, can be efficiently overparametrized, and generalizes well with only polynomially many training points. Thus, the results in Ref. [96] showcase the extreme power of EQNNs and GQML.

B. Implications of our work and future directions

Many concepts and results in our work can be thought of as quantum analogues of existing classical techniques that have enjoyed tremendous success [14]. We envision that GQML will soon be a thriving field, as it provides blueprints to create arbitrary architectures and inductive biases suitable for a given problem. As such, the first direct application of our work is building appropriate schemes to embed classical data into quantum states. Currently, most proposals dealing with classical data use problem-agnostic embedding architectures that completely obviate and destroy the symmetries in the input data [38,113,162]. As such, it is crucial to create embedding schemes that will preserve said symmetries and promote them from the classical to the quantum realm.

As the main challenge of near-term quantum hardware is noise, a most important future research direction is to study the interaction of noise and equivariance. Here, there are two possible paths. On one hand, one can accept that noise will break equivariance and study the effects of such approximate equivariance. Interestingly, it has been observed in the ML literature that mildly breaking equivariance can improve the performance over strictly equivariant models in certain tasks [163]. On the other hand, one can attempt to equivariantize the noise. Being framed in the general superoperator formalism, the present work contains all the necessary tools to study and develop

symmetrization strategies to project noise into the symmetric subspace of a given group representation. Finally, near-term computation will also be limited in resources (e.g., circuit depth, hardware connectivity, etc.), which could prohibit exact equivariance enforcement. Can we derive “cheaper” EQNNs at the cost of only approximate equivariance enforcing? How well do EQNNs perform as a function of the symmetry breaking?

Note added.—There have recently been a number of follow-up works exploring the methods and open questions discussed in our work to develop EQNNs and QML applications in various contexts. For example, Refs. [164–167] use the twirling method to construct QNNs equivariant to finite groups, with applications ranging from calculating molecular force fields and quantum phase detection to image classification. These works all provide numerical results demonstrating improved performance using EQNNs. In Ref. [168], a method is proposed based on spin networks that is shown to be equivalent to the Choi method and it is applied to several lattice Hamiltonian models. In Ref. [169], the authors study the role of choosing representations in EQNN performance. In Ref. [170], the authors study the behavior of EQNNs in the presence of noise and derive strategies to protect equivariance.

ACKNOWLEDGMENTS

This work was partly supported by the U.S. Department of Energy (DOE) through a quantum computing program sponsored by the Los Alamos National Laboratory (LANL) Information Science & Technology Institute. Q.T.N. acknowledges support from the Harvard Quantum Initiative. L.S. was partially supported by the National Science Foundation (NSF) Quantum Leap Challenge Institute (QLCI) for Hybrid Quantum Architectures and Networks (NSF Award No. 2016136). M.R. was partially supported by the NSF through Grants No. DMS-1813149 and No. DMS-2108390. P.J.C. and M.L. were initially supported by the U.S. DOE, Office of Science, Office of Advanced Scientific Computing Research, under the Accelerated Research in Quantum Computing (ARQC) program. P.J.C. and L.S. were also supported by the LANL ASC Beyond Moore’s Law project. F.S. was supported by the Laboratory Directed Research and Development (LDRD) program of LANL under Project No. 20220745ER. M.L. was also supported by the Center for Nonlinear Studies at LANL. M.C. acknowledges support by the LDRD program of LANL under Projects No. 20210116DR and No. 20230049DR.

APPENDIX A: EQUIVARIANCE IN EXISTING QUANTUM ALGORITHMS

In this appendix, we discuss some notable (nonvariational) quantum algorithms from the perspective of group

equivariance. While not explicitly mentioned in the original works, these algorithms rely on equivariant operations. This shows the significance of equivariance in designing quantum algorithms beyond variational QML.

1. Quantum state purity

Computing the purity $\text{Tr}[\rho^2]$ is a canonical task in quantum information theory. Since the purity is unitary invariant, it should be expected that algorithms aiming at measuring it must use unitary-equivariant transformations. We consider two algorithms in Refs. [98,99] and show that, indeed, these use equivariant operations as defined in this work.

The Bell-basis algorithm described in Ref. [98, Fig. 6] starts from two copies of ρ and performs a change of basis to the Bell basis before measuring the observable $CZ^{\otimes n}$ (controlled-phase). This is equivalent to simply measuring the equivariant observable $\text{SWAP}^{\otimes n}$, which belongs to the commutant of the representation $R(U) = U^{\otimes 2}$ of $\mathbb{U}(2^n)$, such that $\text{Tr}[\rho^{\otimes 2} \text{SWAP}^{\otimes n}] = \text{Tr}[U^{\otimes 2} \rho^{\otimes 2} U^{\dagger \otimes 2} \text{SWAP}^{\otimes n}]$ for any $U \in \mathbb{U}(2^n)$.

In contrast, the algorithm in Ref. [99] only uses one copy of ρ and is based off tools from random matrix theory. It starts by appending $m - n$ (where $2^n \ll 2^m$) zero-initialized qubits to the n -qubit state ρ . Then, a random m -qubit unitary is applied on the composite system. Finally, a 2^n -dimensional projective measurement is performed, where one estimates the probability of obtaining some (fixed) outcome $|k\rangle$, $\text{Pr}(k)$. This procedure is repeated many times to estimate $\langle \text{Pr}(k)^2 \rangle$ (with the average taken over the random distribution of m -qubit unitaries), from which one infers the purity as $\text{Tr}[\rho^2] = 2^m \langle \text{Pr}(k)^2 \rangle - 1$. Here, we show that this algorithm is effectively composed of equivariant transformations. Observe that

$$\begin{aligned} \langle \text{Pr}(k)^2 \rangle &= \int_{U \in \mathbb{U}(2^m)} d\mu(U) \\ &\quad \times \text{Tr}[(\rho \otimes |0\rangle \langle 0|_{m-n}) U^\dagger |k\rangle \langle k| U]^2 \\ &= \int_{U \in \mathbb{U}(2^m)} d\mu(U) \\ &\quad \times \text{Tr}[(\rho \otimes |0\rangle \langle 0|_{m-n})^{\otimes 2} U^{\dagger \otimes 2} |kk\rangle \langle kk| U^{\otimes 2}] \\ &= \text{Tr}[(\rho \otimes |0\rangle \langle 0|_{m-n})^{\otimes 2} \\ &\quad \times \int_{U \in \mathbb{U}(2^m)} d\mu(U) U^{\dagger \otimes 2} |kk\rangle \langle kk| U^{\otimes 2}]. \quad (\text{A1}) \end{aligned}$$

We can interpret Eq. (A1) as follows. The first step is $\rho \rightarrow (\rho \otimes |0\rangle \langle 0|)^{\otimes 2}$, which is a $(\mathbb{U}(2^n), R^{\text{in}}, R^{\text{out}})$ -equivariant nonlinear embedding (Definition 9), where we define $R^{\text{in}}(U) = U$ and $R^{\text{out}}(U) = (U \otimes \mathbb{1}_{m-n})^{\otimes 2}$ for $U \in \mathbb{U}(2^n)$. The second step is measuring the observable $\tilde{O} = \int_{U \in \mathbb{U}(2^m)} U^{\dagger \otimes 2} |kk\rangle \langle kk| U^{\otimes 2}$, which commutes with R^{out} due to the invariance of the Haar measure of $\mathbb{U}(2^m)$.

2. Quantum convolutional neural networks (QCNNs)

Quantum convolutional neural networks (QCNNs) have been proposed in Ref. [23]. The architecture presented takes inspiration from classical CNNs [1] and relies on local gates with (potentially) shared parameters in between the gates belonging to the same layer. As an example of application, a QCNN is used to classify phases of the ground states of a Haldane chain [146]. Of particular relevance, the Hamiltonian H_{hal} of the Haldane chain can be verified to commute with the group of symmetry $\mathbb{Z}_2 \times \mathbb{Z}_2 \equiv \{1, r, s, rs\}$, with unitary representation of the generators r and s given by $R(r) = X_{\text{even}} \equiv \prod_{i \text{ even}} X_i$ and $R(s) = X_{\text{odd}} \equiv \prod_{i \text{ odd}} X_i$. Since the nature of the ground state of H_{hal} does not change under the action of these unitaries, we identify $G = \mathbb{Z}_2 \times \mathbb{Z}_2$ as the symmetry group of the task.

For this ground-state classification problem, two QCNNs have been studied in Ref. [23]: one trainable (see Ref. [23, Supplementary Fig. 2]) and one “exact” (see Ref. [23, Fig. 2b]; reproduced in Fig. 13) that is obtained based on the multi-scale entanglement renormalization ansatz (MERA) representation of the ground states of H_{hal} . Remarkably, close inspection of the exact QCNN reveals that it is composed of equivariant layers and measurement. That is, the exact model for this task follows the framework of EQNNs laid down in this work. However, one can see that the choice of trainable model adopted in Ref. [23] does not comply with equivariance requirements, such that we expect that it could be further improved by imposing equivariance. In the following, we briefly discuss

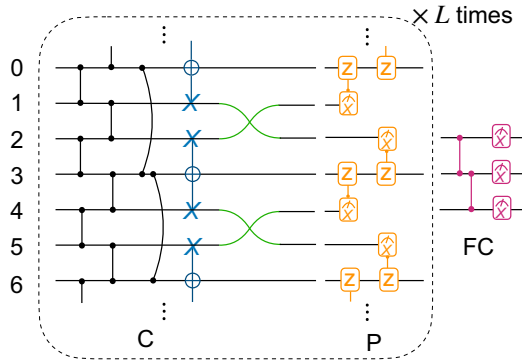


FIG. 13. The reproduction of the “exact” QCNN architectures used in Ref. [23] for the classification of the ground states of the Haldane model: C, convolution; P, pooling; FC, fully connected. The blue three-qubit gates are Toffoli gates, with control qubits taken in the X basis. The orange two-qubit gates apply a Pauli Z onto the target qubit when an X measurement of the control qubit results in an outcome of -1 but leaves the target qubit unchanged otherwise. The black two-qubit gates are CZ gates. The green interleaved lines correspond to a SWAP of two qubits. The C-P structure is repeated L times until the system is left with only three qubits.

how one can identify equivariance of the different layers of the exact QCNN.

First we consider the action of a convolution layer (denoted “C” in Fig. 13 and excluding the SWAP operations) onto $R(r)$. Using the identity

$$\begin{array}{c} \boxed{X} \\ \bullet \\ \text{---} \\ \bullet \\ \text{---} \end{array} = \begin{array}{c} \text{---} \\ \bullet \\ \text{---} \\ \bullet \\ \text{---} \\ \boxed{Z} \end{array}$$

and noting that, due to the connectivity of the controlled- Z (CZ) gates, the resulting Z s unitaries acting on any of the qubits can only be created by pairs (yielding an identity), one can verify that $R(r)$ commutes with the action of all the CZ gates. Additionally, given that both the controls and target of the Toffoli gates act on the eigenbasis of X operators, one can see that $R(r)$ commutes with the action of the whole convolution layer. Similar reasoning shows that such commutativity properties also hold true for $R(s)$ and thus for $R(rs)$. Overall, we find the convolution layers to be (G, R, R) -equivariant.

Second, we consider the action of the pooling layer (denoted P in Fig. 13 and including the SWAP operations). One can verify that $\phi_{\text{pool}} \circ \text{Ad}_{R(g)} = \text{Ad}_{R'(g)} \circ \phi_{\text{pool}}$, where we have denoted as ϕ_{pool} the map realized by the pooling from n to $n/3$ qubits and by $R'(g)$ the representation of $g \in G$ on this reduced space. In particular, R' is defined in equivalence to R but over a reduced number $n/3$ of qubits. Overall, we find the pooling layers to be (G, R, R') -equivariant.

Finally, note that the measurement realized by the fully connected layer (“FC” in Fig. 13) corresponds to a measurement of the Pauli observable $O = ZXZ$ on the three remaining qubits. Notably, $O \in \text{comm}(R^{\text{out}})$, where R^{out} is the representation of G on the remaining qubits and is defined as $R^{\text{out}}(r) = XIX$ and $R^{\text{out}}(s) = IXI$. That is, we find the measurement to be (G, R^{out}) -equivariant, such that the overall exact QCNN follows the requirements of Proposition 1, ensuring invariance of the model.

3. Quantum enhanced experiments

The quantum enhanced experiment in Ref. [21] leverages coherent access to multiple copies of a quantum state obtained from physical experiments to learn its properties. In their “predicting observables” task (Ref. [21, Theorem 1]), the goal is to predict the *absolute value* of an n -qubit Pauli observable O on states of the form $\rho = (\mathbb{1} + 0.9sP)/2^n$, where P is also an n -qubit Pauli string and $s \in \{0, \pm 1\}$. Note that this task is Pauli invariant, since for any Pauli string σ , we have that $|\text{Tr}(\rho O)| = |\text{Tr}(\sigma \rho \sigma O)|$. The first step in the authors’ algorithm is to add a copy: $\rho \rightarrow \rho^{\otimes 2}$, which is a nonlinear equivariant layer as defined in Definition 11. Then, a Bell measurement

is performed, followed by classical postprocessing. Let $G = \{\sigma : n\text{-qubit Pauli strings}\}$ and consider the following representations of G : $R_{\text{def}}(\sigma) = \sigma$ and $R_2(\sigma) = \sigma^{\otimes 2}$. The authors have derived the following equation (see Sec. D.2 of the Supplementary Materials of Ref. [21]):

$$|\text{Tr}[\rho O]| = \mathbb{E} \left[\text{Tr} \left[O \otimes O \bigotimes_{k=1}^n S_k \right] \right], \quad (\text{A2})$$

where S_k is the Bell projector corresponding to the Bell-measurement result on qubit pair k . That is, $S_k = |\Psi_k\rangle \langle \Psi_k|$ for a Bell state $|\Psi_k\rangle$. Expanding Eq. (A2), we have

$$\begin{aligned} |\text{Tr}[\rho O]| &= \sum_{|\Psi_1 \dots \Psi_n\rangle} \text{Tr} [O^{\otimes 2} |\Psi_1 \dots \Psi_n\rangle \langle \Psi_1 \dots \Psi_n| \langle \Psi_1 \dots \Psi_n| \rho^{\otimes 2} |\Psi_1 \dots \Psi_n\rangle] \\ &= \text{Tr} \left[\rho^{\otimes 2} \sum_{|\Psi_1 \dots \Psi_n\rangle} |\Psi_1 \dots \Psi_n\rangle \langle \Psi_1 \dots \Psi_n| O^{\otimes 2} |\Psi_1 \dots \Psi_n\rangle \langle \Psi_1 \dots \Psi_n| \right] \\ &= \text{Tr}[\rho^{\otimes 2} \tilde{O}]. \end{aligned} \quad (\text{A3})$$

We can see that after adding another copy of ρ , the algorithm effectively performs a measurement of the observable \tilde{O} . It is readily verified that \tilde{O} commutes with R_2 using the fact that, for any single-qubit Pauli σ , the operator $\sigma^{\otimes 2}$ admits the Bell states $|\Psi_k\rangle$ as eigenvectors with eigenvalues ± 1 . Hence, the entire algorithm is R_{def} -invariant.

4. Classical shadows

Classical shadows is an efficient protocol for predicting observables on quantum states using randomized measurements [100]. In this protocol, one applies a random unitary U drawn from a unitary ensemble \mathcal{E} on the state ρ , then performs a computational basis measurement to obtain a bit string \mathbf{z} . For example, in Ref. [100], the authors have considered Pauli and Clifford ensembles. Repeating this process many times allows one to predict properties of ρ . We first consider \mathcal{E} to be the n -qubit Clifford group. The expected classical shadow (see Sec. 5.B of the Supplementary Information of Ref. [100]) can be rewritten as

$$\begin{aligned} \mathbb{E}_{\mathbf{z}, U} [(2^n + 1) U^\dagger |\mathbf{z}\rangle \langle \mathbf{z}| U - \mathbb{1}] &= (2^n + 1) \sum_{\mathbf{z}} \frac{1}{|\mathcal{E}|} \sum_{U \in \mathcal{E}} U^\dagger |\mathbf{z}\rangle \langle \mathbf{z}| U \text{Tr} [|\mathbf{z}\rangle \langle \mathbf{z}| U \rho U^\dagger] - \mathbb{1} \\ &= (2^n + 1) \text{Tr}_B \left[\underbrace{\left(\frac{1}{|\mathcal{E}|} \sum_{U \in \mathcal{E}} U^{\dagger \otimes 2} \left(\sum_{\mathbf{z}} |\mathbf{z}\mathbf{z}\rangle \langle \mathbf{z}\mathbf{z}| \right) U^{\otimes 2} \right)}_{\tilde{O}} (\mathbb{1} \otimes \rho) \right] - \mathbb{1}, \end{aligned} \quad (\text{A4})$$

where Tr_B denotes the partial trace over the second subsystem.

Note that the sum over \mathcal{E} is equal to the Haar integral over $\mathbb{U}(2^n)$, as the Clifford group forms a 3-design [171], i.e., $\tilde{O} = \int_{U \in \mathbb{U}(2^n)} d\mu(U) U^{\dagger \otimes 2} O U^{\otimes 2}$, where $O := \sum_{\mathbf{z}} |\mathbf{z}\mathbf{z}\rangle \langle \mathbf{z}\mathbf{z}|$. Thus, \tilde{O} commutes with the representation $R_2(U) := U^{\otimes 2}$ of $\mathbb{U}(2^n)$. We can therefore interpret Eq. (A4) as a composition of three equivariant layers in Definition 9:

$$\rho \xrightarrow{\text{Embedding}} \mathbb{1} \otimes \rho \xrightarrow{\text{Standard}} \tilde{O} \rho \xrightarrow{\text{Pooling}} \text{Tr}_B[\tilde{O}(\mathbb{1} \otimes \rho)]. \quad (\text{A5})$$

The corresponding representations transform as

$$R_{\text{def}} \xrightarrow{(G, R_{\text{def}}, R_2)\text{-equivariant}} R_2 \xrightarrow{(G, R_2, R_2)\text{-equivariant}} R_2 \xrightarrow{(G, R_2, R_{\text{def}})\text{-equivariant}} R_{\text{def}}. \quad (\text{A6})$$

Next, we consider \mathcal{E} to be the Pauli ensemble. The expected classical shadow is

$$\begin{aligned}
& \mathbb{E}_{z,U} \left[\bigotimes_{j=1}^n (3U_j^\dagger |z_j\rangle\langle z_j| U_j - \mathbb{1}) \right] \\
&= \sum_z \frac{1}{|\mathcal{E}|} \sum_{U \in \mathcal{E}} \bigotimes_{j=1}^n (3U_j^\dagger |z_j\rangle\langle z_j| U_j - \mathbb{1}) \text{Tr}[|z\rangle\langle z| U \rho U^\dagger] \\
&= \text{Tr}_B \left[\left(\frac{1}{|\mathcal{E}|} \sum_{U \in \mathcal{E}} \bigotimes_{j=1}^n U_j^{\dagger \otimes 2} \underbrace{\left(\sum_{z_j} 3 |z_j z_j\rangle\langle z_j z_j| - \mathbb{1} \otimes |z_j\rangle\langle z_j| \right)}_{M_j} U_j^{\otimes 2} \right) (\mathbb{1} \otimes \rho) \right] \\
&= \text{Tr}_B \left[\left(\bigotimes_{j=1}^n \int_{U_j \in \text{U}(2)} \underbrace{U_j^{\dagger \otimes 2} M_j U_j^{\otimes 2}}_{\tilde{M}_j} \right) (\mathbb{1} \otimes \rho) \right] \\
&= \text{Tr}_B [\tilde{M}(\mathbb{1} \otimes \rho)], \tag{A7}
\end{aligned}$$

where in the last equality we have used the fact that the one-qubit Pauli group forms a 3-design of $\text{U}(2)$. By explicitly evaluating the integral, we find that $\tilde{M}_j = \text{SWAP}$, thus $\tilde{M} = \text{SWAP}^{\otimes n}$, which commutes with the representation R_2 . A similar composition of equivariant layers to Eq. (A5) thus follows.

APPENDIX B: EQUIVARIANT MAPS AS GENERALIZED GROUP CONVOLUTIONS IN FOURIER SPACE

In this appendix, we provide further details on the interpretation of classical equivariant maps as generalized group convolutions, given in Sec. IV A. More specifically, by looking at the Fourier space, we find that group convolution is a special case of equivariant maps when the representation is the regular representation. We first consider a finite group G and later the Lie group.

Recall that the group convolution of two vectors $a, b \in \mathbb{C}^{|G|}$ (also known as functions mapping group elements to scalars) is defined as

$$(a \otimes b)(u) = \sum_{v \in G} a(uv^{-1})b(v). \tag{B1}$$

Note that the above expression can be rewritten as a matrix vector multiplication as follows:

$$a \otimes b = Ab, \quad \text{where } A = \sum_{u \in G} a(u)R_{\text{left}}(u). \tag{B2}$$

Here, R_{left} denotes the *left regular representation* of the group G , which maps each group element u to a permutation matrix $R_{\text{left}}(u)$ that performs $R_{\text{left}}(u)|v\rangle = |uv\rangle$. For example, in CNNs [1], the convolution matrix A is a circulant matrix, since the group $\mathbb{Z}_n \times \mathbb{Z}_n$ is Abelian.

The key property here is that the group Fourier transform, $F_G := \sum_{v \in G} \sum_{\rho \in \hat{G}} \sqrt{d_\rho/|G|} \sum_{i,j=1}^{d_\rho} \rho(v)_{j,k} |\rho, j, k\rangle\langle v|$, block diagonalizes the left regular representations into irreps as [91]

$$R_{\text{left}}(u) = F_G^\dagger \left(\bigoplus_{\xi \in \hat{G}} \xi(u) \otimes \mathbb{1}_{d_\xi} \right) F_G, \tag{B3}$$

where \hat{G} denotes the set of inequivalent irreps of G and d_ξ is the dimension of the irrep ξ .

Thus, under F_G , the convolution matrix A is block diagonalized as

$$A \cong \bigoplus_{\xi \in \hat{G}} \hat{a}(\xi) \otimes \mathbb{1}_{d_\xi}, \tag{B4}$$

where $\hat{a}(\xi) := \sum_{u \in G} a(u)\xi(u)$ is the *Fourier transform* of the kernel c and \hat{G} is the set of inequivalent irreps.

Comparing Eqs. (13) and (29) to Eq. (B4), we see that equivariant channels generalize group convolution. Group convolutions are then equivariant maps to the regular representation, where the multiplicities m_ξ are equal to the irrep dimensions d_ξ and the basis is the regular Fourier

basis. This is why most of the classical homogeneous ENN architectures [11,12] (ENNs on regular representations) implement equivariant layers via group convolutions.

We complete this appendix with a similar analysis on compact Lie groups. For these groups, an irrep decomposition of the regular representation similar to Eq. (B4) also exists due to the Peter-Weyl theorem [115]. Let $U(L^2(G))$ denote the group of unitary operators on the Hilbert space of L^2 -integrable functions on the group G . Then, the left regular representation, $R_{\text{left}} : G \rightarrow U(L^2(G))$, is defined as $R_{\text{left}}(u)f(x) = f(u^{-1}x)$, for $u \in G$ and $f \in L^2(G)$.

Theorem 6 (Peter-Weyl theorem). Let G be a compact group. Then, the regular representation $R_{\text{left}} : G \rightarrow U(L^2(G))$ is isomorphic to a direct sum of the irreducible representations of G :

$$R_{\text{left}} \cong \bigoplus_{\xi \in \hat{G}} \xi \otimes \mathbb{1}_{d_\xi}. \quad (\text{B5})$$

Above, the isomorphism is the group Fourier transform, which maps functions $f \in L^2(G)$ to operator-valued functions, $\hat{f}(\xi) = \int_G d\mu(u) f(u) \xi(u)$. The Lie group convolution is similarly defined as $(f \otimes k)(u) := \int_G d\mu(v) f(uv^{-1})k(v)$ for $f, k \in L^2(G)$. One can verify that it is possible to derive a block diagonalization of f in the Fourier basis similar to that in Eq. (B4). Note that the sum over inequivalent irreps in Eq. (B5) is infinite, as Lie groups have infinitely many inequivalent irreps. In contrast, for Lie-group representations on finite qubit systems (as we consider in this work), the EQNN equivariant layers only process irreps up to some truncated irreps.

APPENDIX C: DEFERRED PROOFS

In this appendix, we present proofs for some of the results in the main text. For the convenience of the reader, we recall the statement of the theorems and propositions prior to their proofs.

1. Deferred proofs from Sec. IV

Theorem 7 (Free parameters in equivariant channels). Let the irrep decomposition of $R := R^{\text{in}*} \otimes R^{\text{out}}$ be $R(g) \cong \bigoplus_q R_q(g) \otimes \mathbb{1}_{m_q}$. Then, any $(G, R^{\text{in}}, R^{\text{out}})$ -equivariant CPTP channels can be fully parametrized via $\sum_q m_q^2 - C(R^{\text{in}}, R^{\text{out}})$ real scalars, where $C(R^{\text{in}}, R^{\text{out}})$ is a positive constant that depends on the considered representations.

Proof of Theorem 3. Let ϕ be a $(G, R^{\text{in}}, R^{\text{out}})$ -equivariant channel. By Theorem 1 and Lemma 1, the Choi operator J^ϕ is decomposed as $J^\phi \cong \bigoplus_{q=1}^Q \mathbb{1}_{d_q} \otimes J_q^\phi$, where each J_q^ϕ is an operator in an m_q -dimensional subspace corresponding to the irrep decomposition $R := R^{\text{in}*} \otimes R^{\text{out}}$.

For convenience of notation, we will denote as $\mathcal{H}_B \otimes \mathcal{H}_A$ the Hilbert space over which J^ϕ acts. Imposing $J^\phi \geq 0$ (CP) is equivalent to imposing $J_q^\phi \geq 0$ for each irrep q . An m_q -dimensional complex positive semidefinite operator is parametrized by m_q^2 real scalars, for a total of $\sum_q m_q^2$ parameters. Next, we impose TP via $\text{Tr}_B[J^\phi] = \mathbb{1}_A$, where $\mathbb{1}_A$ denotes the identity over \mathcal{H}_A .

Let the change of basis in the irrep decomposition be W , i.e., $J^\phi = W^\dagger \left(\bigoplus_{q=1}^Q \mathbb{1}_{d_q} \otimes J_q^\phi \right) W$, where we have dropped the superscript ϕ for brevity. The TP condition reads

$$\begin{aligned} \text{Tr}_B[J^\phi] &= \mathbb{1}_A \\ &= \sum_j \langle j |_B \otimes \mathbb{1}_A \rangle W^\dagger \left(\bigoplus_{q=1}^Q \mathbb{1}_{d_q} \otimes J_q^\phi \right) W |j\rangle_B \otimes \mathbb{1}_A \\ &= \sum_j T_j^\dagger \left(\bigoplus_{q=1}^Q \mathbb{1}_{d_q} \otimes J_q^\phi \right) T_j, \end{aligned} \quad (\text{C1})$$

where $T_j = W(|j\rangle_B \otimes \mathbb{1}_A)$. Vectorizing the above equation, we can use the property $\text{vec}(M_1 M_2 M_3) = (M_3^\top \otimes M_1) \text{vec}(M_2)$ to obtain

$$\begin{aligned} D \cdot \text{vec} \left(\bigoplus_{q=1}^Q \mathbb{1}_{d_q} \otimes J_q^\phi \right) &= \text{vec}(\mathbb{1}_A), \quad \text{where} \\ D &:= \sum_{j \in \dim(\mathcal{H}_B)} T_j^\top \otimes T_j^\dagger \in \mathbb{C}^{\dim(\mathcal{H}_A)^2 \times (\dim(\mathcal{H}_A) \dim(\mathcal{H}_B))^2}. \end{aligned} \quad (\text{C2})$$

Let \tilde{D} be the $\dim(\mathcal{H}_A)^2 \times \sum_q m_q^2$ matrix the columns of which correspond to the nonzero entries in $\text{vec} \left(\bigoplus_{q=1}^Q \mathbb{1}_{d_q} \otimes J_q^\phi \right)$. Then $\text{rank}(\tilde{D}) = C(R^{\text{in}}, R^{\text{out}})$.

It is readily verified that in the nonequivariant case, i.e., $W = \mathbb{1}$ and J^ϕ , is fully parametrized, the matrix $\tilde{D} = D$ is full row rank, in which case imposing TP reduces $\dim(\mathcal{H}_A)^2$ free parameters as expected. ■

Proposition 3 (Insensitivity to equivalent representations). Consider an EQNN as defined in Definition 8. Then, changing an intermediate representation, R^l , to another representation equivalent to it, $VR^l V^\dagger$, where V is a unitary, does not change the expressibility of the EQNN.

Proof of Proposition 2. Here, we show that changing a representation to another equivalent representation does not change the expressibility of the EQNN. In particular, consider two EQNNs that undergo the same

representations except at one place,

$$\begin{aligned} \mathcal{N} : R^{\text{in}} &\longrightarrow \dots \longrightarrow R^1 \xrightarrow{\mathcal{N}^1} R \xrightarrow{\mathcal{N}^2} R^2 \longrightarrow \dots \\ &\longrightarrow R^{\text{out}}. \end{aligned} \quad (\text{C3})$$

$$\begin{aligned} \mathcal{N}' : R^{\text{in}} &\longrightarrow \dots \longrightarrow R^1 \xrightarrow{\mathcal{N}'^1} R' \xrightarrow{\mathcal{N}'^2} R^2 \longrightarrow \dots \\ &\longrightarrow R^{\text{out}}, \end{aligned} \quad (\text{C4})$$

where $R' = VRV^\dagger$ for some unitary V .

Observe that the set of (G, R_1, R) -equivariant channels is in one-to-one correspondence to the set of (G, R_1, R') -equivariant channels. Indeed, for any (G, R_1, R) -equivariant \mathcal{N}_1 , the channel $\mathcal{N}'_1 = \text{Ad}_V \circ \mathcal{N}_1$ is (G, R_1, R') -equivariant, as we have that

$$\begin{aligned} \mathcal{N}'_1 &= \text{Ad}_R \circ \mathcal{N}^1 \circ \text{Ad}_{R_1^\dagger} \\ &= \text{Ad}_R \circ (\text{Ad}_{V^\dagger} \circ (\mathcal{N}^1)) \circ \text{Ad}_{R_1^\dagger} \Leftrightarrow \mathcal{N}'^1 \\ &= \text{Ad}_{VRV^\dagger} \circ \mathcal{N}^1 \circ \text{Ad}_{R_1^\dagger}. \end{aligned} \quad (\text{C5})$$

Similarly, we obtain $\mathcal{N}'^2 = \mathcal{N}^2 \circ \text{Ad}_{V^\dagger}$. Thus, $\mathcal{N}'^2 \circ \mathcal{N}'^1 = \mathcal{N}^2 \circ \mathcal{N}^1$. A similar argument can be made for changing between equivalent output representations, in which case there is a bijection between the observables that commute with R^{out} and those that commute with $R^{\text{out}'}$. ■

2. Deferred proofs from Sec. V

Theorem 8 (Finite group equivariance). Given a finite group G with generating set S , a linear map ϕ is $(G, R^{\text{in}}, R^{\text{out}})$ -equivariant if and only if

$$\phi \circ \text{Ad}_{R^{\text{in}}(g)} - \text{Ad}_{R^{\text{out}}(g)} \circ \phi = 0, \quad \forall g \in S. \quad (\text{C6})$$

Proof of Theorem 4. Given a finite group G and a generating set $S = \{h_1, \dots, h_{|S|}\} \subset G$, we can identify any group element g with a sequence $\kappa = (\kappa_1, \kappa_2, \dots, \kappa_N)$ where $\kappa_i \in \{1, \dots, |S|\}$, such that $g = \prod_{i=1}^N h_{\kappa_i}$ [60]. Assuming that the equivariance condition is satisfied for the generating set, i.e., for any $h \in S$,

$$\text{Ad}_{R^{\text{out}}(h)} \circ \phi \circ \text{Ad}_{R^{\text{in}}(h)}^\dagger = \phi. \quad (\text{C7})$$

Then, we can readily show that it is also satisfied for any $g \in G$:

$$\begin{aligned} &\left(\text{Ad}_{R^{\text{out}}(h_{\kappa_1})} \circ \dots \circ \text{Ad}_{R^{\text{out}}(h_{\kappa_1})} \right) \\ &\times \phi \circ \left(\text{Ad}_{R^{\text{in}}(h_{\kappa_1})} \circ \dots \circ \text{Ad}_{R^{\text{in}}(h_{\kappa_1})} \right)^\dagger = \phi, \end{aligned} \quad (\text{C8})$$

where we have applied Eq. (C7) N times. ■

Theorem 9 (Lie-group equivariance). Given a compact Lie group G with a Lie algebra \mathfrak{g} generated by s such that exponentiation is surjective, a linear map ϕ is $(G, R^{\text{in}}, R^{\text{out}})$ -equivariant if and only if

$$\text{ad}_{r^{\text{out}}(a)} \circ \phi - \phi \circ \text{ad}_{r^{\text{in}}(a)} = 0, \quad \forall a \in s, \quad (\text{C9})$$

where r^{in} and r^{out} are the representations of G induced by R^{in} and R^{out} .

Proof of Theorem 5. Let us start by recalling that for any element $a \in \mathfrak{g}$, there is a corresponding $e^a \in G$. If ϕ is $(G, R^{\text{in}}, R^{\text{out}})$ -equivariant, then $\text{Ad}_{R^{\text{out}}(g)} \circ \phi = \phi \circ \text{Ad}_{R^{\text{in}}(g)}$, $\forall g \in G$. Differentiating this expression yields that $\text{ad}_{r^{\text{out}}(a)} \circ \phi = \phi \circ \text{ad}_{r^{\text{in}}(a)}$, where $g = e^a$. Since this holds for any element a in \mathfrak{g} , it must hold for the generating set of \mathfrak{g} .

We now prove the other direction by first showing that $\text{ad}_{r^{\text{out}}(a)} \circ \phi = \phi \circ \text{ad}_{r^{\text{in}}(a)}$, $\forall a \in \mathfrak{g}$ assuming that this relation holds for the generating set $s = \{a_1, \dots, a_{|s|}\}$ of \mathfrak{g} . First, we consider the element $[a_i, a_j]$ in \mathfrak{g} as follows:

$$\begin{aligned} &\text{ad}_{r^{\text{out}}([a_i, a_j])} \circ \phi \\ &= \text{ad}_{r^{\text{out}}(a_i)} \circ \text{ad}_{r^{\text{out}}(a_j)} \circ \phi - \text{ad}_{r^{\text{out}}(a_j)} \circ \text{ad}_{r^{\text{out}}(a_i)} \circ \phi \\ &= \phi \circ \text{ad}_{r^{\text{in}}(a_i)} \circ \text{ad}_{r^{\text{in}}(a_j)} - \phi \circ \text{ad}_{r^{\text{in}}(a_j)} \circ \text{ad}_{r^{\text{in}}(a_i)} \\ &= \phi \circ \text{ad}_{r^{\text{in}}([a_i, a_j])}. \end{aligned}$$

Recursively applying the above calculation, we find that $\text{ad}_{r^{\text{out}}(a')} \circ \phi = \phi \circ \text{ad}_{r^{\text{in}}(a')}$ for any nested commutator $a' = [a_{i_1}, [a_{i_2}, \dots]]$. Hence, $\text{ad}_{r^{\text{out}}(a)} \circ \phi = \phi \circ \text{ad}_{r^{\text{in}}(a)}$, $\forall a \in \mathfrak{g}$, since s generates \mathfrak{g} . It follows that $e^{\text{ad}_{r^{\text{out}}(a)}} \circ \phi = \phi \circ e^{\text{ad}_{r^{\text{in}}(a)}}$. Since exponentiation is surjective, for all $g \in G$ there is a corresponding $a \in \mathfrak{g}$ such that $g = e^a$ and accordingly $e^{\text{ad}_{r(a)}} = \text{Ad}_{R(g)}$. Therefore, $\text{Ad}_{R^{\text{out}}(g)} \circ \phi = \phi \circ \text{Ad}_{R^{\text{in}}(g)}$, $\forall g \in G$. ■

3. Relaxing the assumption of surjectivity

In Theorem 5, we have assumed that the exponentiation map $e : \mathfrak{g} \rightarrow G$ between the Lie algebra and group is surjective. This lets us work interchangeably at both algebra or group level. However, even when the surjectivity assumption is relaxed (i.e., when the exponential of the algebra generates only some subset of the Lie group), there may still be a finite set of generators such that equivariance with respect to this set implies equivariance with the entire group. This happens since the exponential map takes a Lie algebra to a connected component of the Lie group (that containing the identity) but Lie groups can have multiple connected components. In cases in which there are finitely many connected components and the quotient of the Lie group with a finite subgroup H yields a connected Lie

group G/H , the theory does not require much adjustment. When this holds, one can write the group as

$$G = e^{\mathfrak{g}}H. \quad (\text{C10})$$

That is, for all $g \in G$, there exists $a \in \mathfrak{g}$ and $h \in H$ such that $g = e^a h$. In Ref. [60], the authors achieve Lie-group equivariance by imposing the constraint over a basis of \mathfrak{g} and a generating set for H . Here, we leverage Theorem 5 to reduce the number of constraints by imposing them over the connected component generated by \mathfrak{g} to a generating set of it.

Theorem 10. Consider a Lie group G with a Lie algebra \mathfrak{g} generated by $s = \{a_1, \dots, a_{|S|}\}$ such that $G = e^{\mathfrak{g}}H$, where H is a subgroup generated by $S = \{h_1, \dots, h_{|S|}\}$. Then, a linear map ϕ is $(G, R^{\text{in}}, R^{\text{out}})$ -equivariant if and only if

$$\text{ad}_{R^{\text{out}}(a_i)} \circ \phi = \phi \circ \text{ad}_{R^{\text{in}}(a_i)} \quad \forall a_i \in s, \quad (\text{C11})$$

$$\text{Ad}_{R^{\text{out}}(h_i)} \circ \phi = \phi \circ \text{Ad}_{R^{\text{in}}(h_i)}, \quad \forall h_i \in S. \quad (\text{C12})$$

Proof. As shown in the proof of Theorem 5, if ϕ is equivariant, then it commutes with the representation on the algebra level. Further, it also clearly commutes with the elements of H . The other direction also follows similarly from the proof of Theorem 5. From there, we have that $\text{Ad}_{R^{\text{out}}(g)} \circ \phi = \phi \circ \text{Ad}_{R^{\text{in}}(g)}$, $\forall g \in e^{\mathfrak{g}}$. Then, by the assumption that ϕ commutes on H and $G = e^{\mathfrak{g}}H$, we have that, for any $g = e^a h$,

$$\begin{aligned} \text{Ad}_{R^{\text{out}}(g)} \circ \phi &= \text{Ad}_{R^{\text{out}}(e^a h)} \circ \phi \\ &= \text{Ad}_{R^{\text{out}}(e^a)} \circ \text{Ad}_{R^{\text{out}}(h)} \circ \phi \\ &= \phi \circ \text{Ad}_{R^{\text{in}}(e^a)} \circ \text{Ad}_{R^{\text{in}}(h)} \\ &= \phi \circ \text{Ad}_{R^{\text{in}}(g)}. \quad \blacksquare \end{aligned}$$

As an example, consider $\mathbb{S}\mathbb{O}(3)$ and $\mathbb{O}(3)$. Exponentiation of $\mathfrak{so}(3)$ is surjective on $\mathbb{S}\mathbb{O}(3)$ and thus we can take $H = \{\mathbb{1}_3\}$ (the trivial group). However, $\mathbb{O}(3)$ consists of two connected components corresponding to ± 1 determinant. Due to determinant being multilinear, one can map between these components with the subgroup $H = \{\mathbb{1}_3, (-1) \oplus \mathbb{1}_2\}$.

APPENDIX D: ADVANCED TWIRLING METHODS

In the main text, we have discussed how to obtain equivariant maps via twirling. Here, we discuss several advanced techniques to implement the twirling operator. First, we will show how the twirl of an operator or a map can be obtained via the Weingarten calculus. Next, we will showcase two methods for in-circuit twirling, and one for approximate twirling.

1. Weingarten calculus

The Weingarten calculus [127,128,172] is an extremely powerful tool that can be used to find the exact expression of the twirl of an operator over a group. At its core, the Weingarten calculus leverages the key property that twirling is equivalent to projecting into the commutant. Hence, if the commutant of the representation is well known (e.g., via the Schur-Weyl duality [79,145,173]), then one can analytically find an expression for the twirled operator in terms of its components over a basis of the commutant. In what follows, we exemplify the Weingarten calculus for twirling an operator when $R^{\text{in}} = R^{\text{out}} = R$. For a channel, one can use the techniques presented here by finding the components of the Choi operator in the commutant of the $R^{\text{in}*} \otimes R^{\text{out}}$, which can be identified using the *mixed* Schur-Weyl duality [79,173].

Given an operator X , a group G , and a representation R , we already know that the twirl is a projection over the commutant; i.e.,

$$\begin{aligned} \mathcal{T}_G[X] &= \int_G d\mu(g) R(g) X R(g)^\dagger \\ &= \sum_{i=1}^{\dim(\text{comm}(R))} c_\mu(X) P_i, \quad \text{with } P_i \in \text{basis}(\text{comm}(R)). \end{aligned} \quad (\text{D1})$$

Hence, in order to solve Eq. (D1), one needs to determine the $\dim(\text{comm}(R))$ unknown coefficients $\{c_i\}_{i=1}^{\dim(\text{comm}(R))}$. The foregoing can be achieved by finding $\dim(\text{comm}(R))$ such equations and solving a linear system problem. In particular, note that changing $X \rightarrow X P_j$ for some $P_j \in \text{basis}(\text{comm}(R))$ leads to

$$\begin{aligned} \mathcal{T}_G[X P_j] &= \int_G d\mu(g) R(g) X P_j R(g)^\dagger \\ &= \int_G d\mu(g) R(g) X R(g)^\dagger P_j \\ &= \sum_{i=1}^{\dim(\text{comm}(R))} c_i(X) P_i P_j, \end{aligned} \quad (\text{D2})$$

where we have used the fact that P_j commutes with all $R(g)$. Then, taking the trace on both sides leads to

$$\text{Tr}[X P_j] = \text{Tr}[\mathcal{T}_G[X P_j]] = \sum_{i=1}^{\dim(\text{comm}(R))} c_i(X) \text{Tr}[P_i P_j]. \quad (\text{D3})$$

Repeating Eq. (D3) for all P_j 's in $\text{basis}(\text{comm}(R))$ leads to $\dim(\text{comm}(R))$ equations. Thus, one can find the vector of unknown coefficients $\mathbf{c}(X) = (c_1(X), \dots, c_{\dim(\text{comm}(R))}(X))$ by solving $A \cdot \mathbf{c}(X) = \mathbf{b}(X)$, where $\mathbf{b}(X) =$


$(\text{Tr}[XP_1], \dots, \text{Tr}[XP_{\dim(\text{comm}(R))}])$. Here, A is the so-called Gram matrix, a $\dim(\text{comm}(R)) \times \dim(\text{comm}(R))$ symmetric matrix with entries $[A]_{ij} = \text{Tr}[P_i P_j]$. One can then solve the linear-system problem by inverting the Gram matrix, as $\mathbf{c}(X) = A^{-1} \cdot \mathbf{b}(X)$. The matrix A^{-1} is known as the Weingarten matrix.

2. In-circuit twirling with ancillas or classical randomness

While we usually find the set of equivariant channels analytically, in the cases of small finite groups, we note that one could perform the twirling directly on the quantum circuit using the following unitaries:


$$\begin{aligned} \mathbf{U}^{\text{in}} &= \sum_{g \in G} |g\rangle \langle g| \otimes R^{\text{in}}(g)^\dagger, \\ \mathbf{U}^{\text{out}} &= \sum_{g \in G} |g\rangle \langle g| \otimes R^{\text{out}}(g), \end{aligned} \quad (\text{D4})$$

along with $a = \log_2 |G|$ ancilla qubits initialized to the uniform superposition state $H^{\otimes a} |0\rangle$. That is, the in-circuit twirling of a n -to- m -qubit channel ϕ can be realized via the following circuit:



It can be readily verified that this circuit performs the twirling formula in Eq. (25). With this, ϕ can be any parametrized channel native to the circuit platform.

Alternatively, the ancilla qubits can be replaced by classical randomness. That is, we classically sample a group element g and then apply $R^{\text{in}\dagger}(g)$ and $R^{\text{out}}(g)$ as follows:



The latter method can be favorable on near-term devices. Furthermore, the Hoeffding's bound implies that only

$O(\log |G|)$ classical samples are needed to achieve a good approximation of the twirled channel.

One disadvantage of in-circuit twirling, however, is that albeit ensuring equivariance, we lose the parameter-count reduction, as in the case in which we first compute equivariant channels analytically before parametrization.

3. Recursive approximate twirling

For Lie groups with more intricate representation theory, computing the twirling formula can quickly become complex and difficult. Instead, In Ref. [129], the authors have provided an algorithm for approximating twirling operators that converges *exponentially* fast in the number of Haar-random samples of group elements. While the authors have not mentioned this, their proofs do not rest upon any assumptions beyond that the representations that they consider are unitary representations of compact Lie groups and the input of the twirling formula is self-adjoint. Their algorithm can be applied to our case when one can efficiently sample from the Haar measure and is summarized in Algorithm 1. This approximate twirling algorithm can also be implemented in circuit using classical randomness, similarly to what we saw earlier in the case of finite groups.

APPENDIX E: IMPLEMENTING AND OPTIMIZING EQUIVARIANT CHANNELS

1. Channel compiling

In the process of creating equivariant QNNs, we consider not just equivariant unitaries but also more general quantum channels. Via the Stinespring dilation theorem, any channel $\phi : \mathcal{B}(\mathcal{H}_A) \rightarrow \mathcal{B}(\mathcal{H}_B)$ can be represented as a unitary operation on a larger space:

$$\phi(\rho) = \text{Tr}_E[U(\rho \otimes |e\rangle \langle e|)U^\dagger], \quad (\text{E1})$$

where $|e\rangle$ is a fixed reference state on the environment E . The size of this environment system is directly related to the Kraus rank of ϕ . Recall that a quantum channel can be written as

$$\phi(\rho) = \sum_i K_i \rho K_i^\dagger, \quad (\text{E2})$$

where we say that $\{K_i\}$ are the Kraus operators of the channel. Note that the spectral decomposition of the Choi

ALGORITHM 1. Recursive approximate twirling [129].

Require: Group G , unitary representations $R^{\text{in}}, R^{\text{out}}$, map ϕ , tolerance $\varepsilon > 0$, $\delta > 0$.

Ensure: Output map $\tilde{\phi}$ such that $\|\tilde{\phi} - \mathcal{T}_G[\phi]\|_{HS} < \varepsilon$ with probability at least $1 - \delta$.

▷ HS: Hilbert-Schmidt norm

$\tilde{\phi} \leftarrow \phi$

for $k \leftarrow 1$ to $K = O(\log \frac{1}{\delta\varepsilon})$ **do**

 Draw random $g \in G$

$\tilde{\phi} \leftarrow \frac{1}{2} (\tilde{\phi} + R^{\text{out}}(g) \circ \tilde{\phi} \circ R^{\text{in}\dagger}(g))$

operator yields one possible set of Kraus operators. At most $\dim(\mathcal{H}_A)\dim(\mathcal{H}_B)$ Kraus operators are necessary to represent any channel [116]. One can then define a unitary with action,

$$U^\phi = \sum_i K_i \otimes |i\rangle\langle e| + U', \quad (\text{E3})$$

where U' is some arbitrary operator that completes U to be unitary. Thus, to represent a channel with k Kraus operators, an environment of dimension at least k suffices. Working with qubit systems, if the input space is n qubits and the output m , then the maximum Kraus rank is 2^{m+n} , which may require an environment of up to $m+n$ qubits.

We will not go into great detail on how to perform this circuit compilation but, rather, refer the interested reader to the considerable body of work on compilation. For example, a software package for this decomposition with a nearly optimal controlled-NOT (CNOT) count can be found in Ref. [174], with corresponding theory in Refs. [175, 176]. For more general works on circuit compilation, we direct the reader to [25,98,134–136,177,178].

2. Converting vectorized channels to Choi operators

In solving the null space for equivariant maps, we work with vectorized channels; i.e., $\phi \mapsto \bar{\phi} = \sum_{i,j} \phi_{i,j} |P_i\rangle\rangle\langle\langle P_j|$, where P_j and P_i are Pauli strings. In some references, this is referred to as a transfer matrix. The canonical construction of transfer matrices is as follows. Consider the Kraus-operator form of a channel, $\phi(\rho) = \sum_i K_i \rho K_i^\dagger$. The vectorization map $\text{vec}(X)$ takes a matrix X to a vector through the mapping

$$\text{vec}(|i\rangle\langle j|) = |i\rangle \otimes |j\rangle, \quad (\text{E4})$$

with linear extension. Then, using the identity $\text{vec}(AXB) = (B^T \otimes A)\text{vec}(X)$, one can write a matrix representation of ϕ as

$$\bar{\phi} = \sum_i K_i^* \otimes K_i. \quad (\text{E5})$$

This is the transfer matrix. Further, the transfer matrix can be directly mapped to the Choi operator via

$$\tau^{\mathcal{N}} = N^\Gamma, \quad (\text{E6})$$

where Γ is an involution map such that

$$\langle i,j | \bar{\phi} |k,l\rangle = \langle l,j | \bar{\phi}^\Gamma |k,i\rangle. \quad (\text{E7})$$

For us to apply this identity, we need to convert from the transfer matrix in terms of Pauli strings to that of the computational basis. As any $|i\rangle\langle j|$ can be written as a sum

over Pauli strings, there is some change of basis U from Pauli strings to the computational operator basis. Explicitly, the columns of U will be vectors corresponding to the expansion of Pauli strings in the computational basis. If we obtain a matrix in the Pauli basis, X , we can then write it in the computational basis via UAV^{-1} , where U is the change of basis on the output space and V is that on the input space. Further, if we take Pauli strings to be normalized such that $\text{Tr}[P_i P_j] = \delta_{i,j}$, then U and V can be assumed to be unitary matrices. In this case,

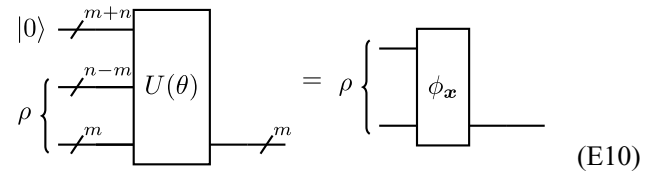
$$J^\phi = (U\bar{\phi}V^\dagger)^\Gamma. \quad (\text{E8})$$

3. Optimizing equivariant channels

We now provide a strategy for optimizing n -to- m -qubit equivariant channels. Assume that we have found a basis $\{\phi_i\}$ for such channels via methods outlined in this work. Equivariant channels in the span of this basis can then be written as

$$\phi_x[\cdot] = \sum_{i=1} x_i \phi_i[\cdot] + \frac{\text{Tr}[\cdot]}{2^m} \mathbb{1}. \quad (\text{E9})$$

The coefficients must be constrained so that ϕ_x is CPTP. For convenience, we have fixed $\phi(\rho) = (\text{Tr}[\rho]/2^m)\mathbb{1}$ as one of the basis elements. Without loss of generality, we can then take all ϕ_i to be trace annihilating, i.e., $\text{Tr}[\phi_i(\rho)] = 0$. Note that depending on the methods used to find these maps, we can bake in CP or TP or even both. Given a valid set of parameters for this pooling layer, we can *classically* solve the following circuit-compilation problem as described before:



The diagram shows a quantum circuit. On the left, there are two input registers: a register of $m+n$ qubits initialized to $|0\rangle$ and a register of m qubits initialized to ρ . These two registers are combined and pass through a unitary gate $U(\theta)$. The output of $U(\theta)$ is a state ρ on m qubits and a state ϕ_x on $m+n$ qubits. The equation is labeled (E10).

The top $2n$ qubits are discarded. $U(\theta)$ is a general $(2n+m)$ -qubit unitary. Note that $m+n$ is an upper bound on the number of ancilla qubits needed but depending on the ranks of the basis channels, we could potentially need fewer qubits. After solving this classical circuit-compilation problem, one can then implement it on a real quantum circuit.

Now that we have a way to implement and parametrize equivariant channels, we can train them via two approaches:

- (i) projected gradient descent (GD) in the circuit parameters θ space (parameter-shift rule works)
- (ii) projected GD in the classical variables x space

For more details on projected GD, we refer the reader to Ref. [126]. Other constrained optimizers can be used as well.

In the first approach, we treat x_i as functions of θ and train θ via the projected GD algorithm, i.e., perform the following in each training step:

$$\begin{aligned} \theta^{k+\frac{1}{2}} &= \theta_k - \eta \nabla_{\theta} \mathcal{L} \quad (\text{regular GD}), \\ \theta^{k+1} &= \min_{\theta'} \|\theta' - \theta^{k+\frac{1}{2}}\| \quad (\text{projection}), \\ &\text{subject to: } \phi_{x(\theta')} \text{ is CPTP.} \end{aligned} \quad (\text{E11})$$

In the second approach, we treat θ as functions of x_i and train x_i . In this case, however, we cannot use finite-difference methods to compute the derivatives with respect to these parameters:

$$\begin{aligned} x^{k+\frac{1}{2}} &= x^k - \eta \nabla_x \mathcal{L} \quad (\text{regular GD}), \\ x^{k+1} &= \min_{x'} \|x' - x^{k+\frac{1}{2}}\| \quad (\text{projection}), \\ &\text{subject to: } \phi_x \text{ is CPTP.} \end{aligned} \quad (\text{E12})$$

One might also want to recycle the ancilla qubits. This is possible if we can replace the partial trace operation by a measurement on the ancillas followed by a controlled unitary (on the possible outcomes). This requires the channel to be unital, which is the case if the output representation is irreducible.

Lemma 3. Equivariant channels the output representation of which is irreducible are unital.

Proof. We require that $\varphi(R^{\text{in}}(g)\rho R^{\text{in}}(g)^{\dagger}) = R^{\text{out}}(g)\varphi(\rho)R^{\text{out}}(g)^{\dagger}$ for any $g \in G$. Substituting $\rho = \mathbb{1}^{\otimes n}$ and applying Schur's lemma, we find that $\varphi(\mathbb{1}^{\otimes n}) = \mathbb{1}^{\otimes m}$. ■

APPENDIX F: SU(2)-EQUIVARIANT 2-TO-1-QUBIT AND 1-TO-2-QUBIT CHANNELS

1. From equivariant maps to channels

Let us first define the five ($\text{SU}(2)$, $U^{\otimes 2}$, U)-equivariant linear maps, which we denote as

$$\begin{aligned} \phi_1(\rho) &= \text{Tr}[\rho] \frac{\mathbb{1}}{2}, & \phi_2(\rho) &= \text{Tr}[\rho \text{ SWAP}] \frac{\mathbb{1}}{2}, \\ \phi_3(\rho) &= \text{Tr}_A[\rho], \\ \phi_4(\rho) &= \text{Tr}_B[\rho], & \phi_5(\rho) &= \sum_{ijk=1}^3 \text{Tr}[\rho \sigma_i \sigma_j] \epsilon_{ijk} \sigma_k. \end{aligned} \quad (\text{F1})$$

We will refer to ϕ_5 as the *cross-product* channel. Here, we will give slightly more detail on how we find the feasible region for channels. To see that ϕ_2 may act nontrivially

on the trace of an input, we consider its vectorization in the Pauli basis. The row corresponding to $|\mathbb{1}_{\text{out}}\rangle\langle P|$ contains all information about $\text{Tr}[\phi_2(\rho)]$, as $\mathbb{1}$ is the only Pauli string of nonzero trace. One can show that $\text{Tr}[\text{SWAP}(\sigma_i \otimes \sigma_j)] = 2\delta_{ij}$ and thus $\phi_2(\sigma_i \otimes \sigma_i) = \mathbb{1}/2$. But this is problematic, as $\text{Tr}[\phi_j(\sigma_i \otimes \sigma_i)] = 0$ for $j \in \{1, 3, 4, 5\}$. Thus, ϕ_2 increases the trace of a state such as $\frac{1}{2}(I + XX)$ but none of the other maps could cancel out this increase. Thus, we drop ϕ_2 from our set of maps.

To finish finding the feasible region, we will make some modifications to our basis elements. That is, we want to modify the basis set such that all elements except for ϕ_1 are trace annihilating. Then we can, without loss of generality, fix the coefficient for the trace-preserving channel to be 1. As the cross-product channel is traceless, we need only modify the partial trace channels. This is easy in the Pauli string basis: simply remove the entry in the upper left-hand corner corresponding to $\mathbb{1} \otimes \mathbb{1} \mapsto \mathbb{1}$. For example, the Tr_A channel becomes

$$\begin{pmatrix} 0 & 0 & 0 & 0 & 0 & \dots \\ 0 & 2 & 0 & 0 & 0 & \dots \\ 0 & 0 & 2 & 0 & 0 & \dots \\ 0 & 0 & 0 & 2 & 0 & \dots \end{pmatrix}. \quad (\text{F2})$$

With these modified channels, we know that the set of equivariant channel can be characterized as

$$\{x, y, z \in \mathbb{R}^3 : J^{\phi_1} + xJ^{\phi_5} + yJ^{\phi_3} + zJ^{\phi_4} \geq 0\}. \quad (\text{F3})$$

Requiring the eigenvalues of this linear combination to be non-negative yields the feasible region,

$$\begin{aligned} \{x, y, z : y + z \leq 1 \text{ and } y + z \\ \geq \sqrt{3x^2 + 4(y^2 - yz + z^2)} - 1\}. \end{aligned} \quad (\text{F4})$$

2. Action of SU(2)-equivariant maps

Here, we further analyze the action of the 2-to-1-qubit maps. This analysis will show that different channels can “see” different parts of the input state and hence that they are complementary. First, define the Bell-basis states as

$$\begin{aligned} |\beta_{00}\rangle &= \frac{1}{\sqrt{2}}(|00\rangle + |11\rangle), & |\beta_{01}\rangle &= \frac{1}{\sqrt{2}}(|00\rangle - |11\rangle), \\ |\beta_{10}\rangle &= \frac{1}{\sqrt{2}}(|01\rangle + |10\rangle), & |\beta_{11}\rangle &= \frac{1}{\sqrt{2}}(|01\rangle - |10\rangle). \end{aligned} \quad (\text{F5})$$

Then, consider a two-qubit quantum state in the Bell basis:

$$\rho = \begin{pmatrix} a_{11} & a_{12} & a_{13} & a_{14} \\ a_{12}^* & a_{22} & a_{23} & a_{24} \\ a_{13}^* & a_{23}^* & a_{33} & a_{34} \\ a_{14}^* & a_{24}^* & a_{34}^* & a_{44} \end{pmatrix}, \quad (\text{F6})$$

where $a_{11} + a_{22} + a_{33} + a_{44} = 1$.

One can readily find that

$$\phi_1(\rho) = \frac{1}{2} \begin{pmatrix} a_{11} + a_{22} + a_{33} + a_{44} & 0 \\ 0 & a_{11} + a_{22} + a_{33} + a_{44} \end{pmatrix}, \quad (\text{F7})$$

$$\phi_2(\rho) = \frac{1}{2} \begin{pmatrix} a_{11} + a_{22} + a_{33} - a_{44} & 0 \\ 0 & a_{11} + a_{22} + a_{33} - a_{44} \end{pmatrix}, \quad (\text{F8})$$

$$\phi_3(\rho) = \begin{pmatrix} \frac{1}{2} + \text{Re}[a_{12} - a_{34}] & \text{Re}[a_{13} + a_{24}] + i\text{Im}[a_{23} + a_{14}] \\ \text{Re}[a_{13} + a_{24}] - i\text{Im}[a_{23} + a_{14}] & \frac{1}{2} - \text{Re}[a_{12} - a_{34}] \end{pmatrix}, \quad (\text{F9})$$

$$\phi_4(\rho) = \begin{pmatrix} \frac{1}{2} + \text{Re}[a_{12} + a_{34}] & \text{Re}[a_{13} - a_{24}] + i\text{Im}[a_{23} - a_{14}] \\ \text{Re}[a_{13} - a_{24}] - i\text{Im}[a_{23} - a_{14}] & \frac{1}{2} - \text{Re}[a_{12} + a_{34}] \end{pmatrix}, \quad (\text{F10})$$

$$\phi_5(\rho) = 4 \begin{pmatrix} -\text{Im}[a_{34}] & \text{Im}[a_{24}] - i\text{Re}[a_{14}] \\ \text{Im}[a_{24}] + i\text{Re}[a_{14}] & \text{Im}[a_{34}] \end{pmatrix}. \quad (\text{F11})$$

These equations show how different channels combine different pieces of the information of ρ .

3. Cross-product channel

We now take a closer look at ϕ_5 , which we call the *cross-product* channel $\mathcal{CP} : \mathcal{H}^{\otimes 2} \rightarrow \mathcal{H}$, where \mathcal{H} is the single-qubit Hilbert space. The action of the \mathcal{CP} (not to be confused with complete positivity) channel is as follows:

$$\begin{aligned} \mathcal{CP}(\rho) &= \sum_{ijk=1}^3 \text{Tr}[\rho \sigma_i \sigma_j] \epsilon_{ijk} \sigma_k \\ &= \text{Tr}[\rho(YZ - ZY)]X + \text{Tr}[\rho(ZX - XZ)]Y \\ &\quad + \text{Tr}[\rho(XY - YX)]Z, \end{aligned} \quad (\text{F12})$$

where $\sigma_\mu \in \{X, Y, Z\}$ for $\mu = i, j, k$.

First, note that in the above form, the \mathcal{CP} channel is not truly a channel, as it is neither trace preserving nor completely positive. Rather, $\text{Tr}[\mathcal{CP}(\rho)] = 0$. To be a channel, we must consider superoperators of the form $\phi + \alpha\mathcal{CP}$, where ϕ is some trace-preserving map. For simplicity, we consider $\phi(\rho) = (\text{Tr}[\rho]/2)\mathbb{1}$. By solving for the eigenvalues of the Choi operator of $\phi + \alpha\mathcal{CP}$, one can show that this is a channel for $\alpha \in [-\frac{1}{\sqrt{3}}, \frac{1}{\sqrt{3}}]$.

Now that we know when this can actually be physical, we would like to better understand the action of the \mathcal{CP} channel. As before, we recall the four Bell-basis states:

$$\begin{aligned} |\beta_{00}\rangle &= \frac{1}{\sqrt{2}}(|00\rangle + |11\rangle), & |\beta_{01}\rangle &= \frac{1}{\sqrt{2}}(|00\rangle - |11\rangle), \\ |\beta_{10}\rangle &= \frac{1}{\sqrt{2}}(|01\rangle + |10\rangle), & |\beta_{11}\rangle &= \frac{1}{\sqrt{2}}(|01\rangle - |10\rangle), \end{aligned} \quad (\text{F13})$$

where $|\beta_{00}\rangle$, $|\beta_{01}\rangle$, and $|\beta_{10}\rangle$ are eigenstates of the SWAP operator with eigenvalue $+1$, while $|\beta_{11}\rangle$ is an eigenvalue of the SWAP operator with eigenvalue -1 . Then, we note the following operator expansion in the Bell basis:

$$YZ - ZY = 2i(|\beta_{01}\rangle\langle\beta_{11}| - |\beta_{11}\rangle\langle\beta_{01}|), \quad (\text{F14})$$

$$ZX - XZ = 2(|\beta_{00}\rangle\langle\beta_{11}| + |\beta_{11}\rangle\langle\beta_{00}|), \quad (\text{F15})$$

$$XY - YX = -2i(|\beta_{10}\rangle\langle\beta_{11}| - |\beta_{11}\rangle\langle\beta_{10}|). \quad (\text{F16})$$

We can then express the density matrix in the Bell basis (ordered as $\{|\beta_{00}\rangle, |\beta_{01}\rangle, |\beta_{10}\rangle, |\beta_{11}\rangle\}$),

$$\rho = \begin{pmatrix} \begin{matrix} \square & & & \\ & \square & & \\ & & \square & \\ & & & \square \end{matrix} & \begin{matrix} a_{14} \\ a_{24} \\ a_{34} \end{matrix} \\ \begin{matrix} a_{14}^* \\ a_{24}^* \\ a_{34}^* \end{matrix} & \square \end{pmatrix}, \quad (\text{F17})$$

where the 3×3 and 1×1 diagonal blocks correspond to the symmetric and antisymmetric subspaces, respectively. Then, the matrix elements a_{14} , a_{24} , and a_{34} determine if the state is in a superposition of symmetric and antisymmetric Bell-basis states.

With the foregoing, one can verify that

$$\text{Tr}[\rho(YZ - ZY)] = 4\text{Im}[a_{24}],$$

$$\text{Tr}[\rho(ZX - XZ)] = 4\text{Re}[a_{14}],$$

$$\text{Tr}[\rho(XY - YX)] = -4\text{Im}[a_{34}].$$

Note that the action of the \mathcal{CP} channel is to check if ρ is in a superposition of states with different symmetries.

As such, at the output of the map, the coefficients associated with the different Pauli operators correspond to the entries of the matrix entries that account for superposition between the antisymmetric state and the three different symmetric states. Hence, the \mathcal{CP} channel outputs the zero matrix for any state that is block diagonal in the symmetric and antisymmetric subspaces (such as $\rho = \sigma^{\otimes 2}$ for any single-qubit state σ).

Here, let us note an interesting fact, namely, that the \mathcal{CP} channel, in its vanilla version, has an asymmetry embedded into it: it only accounts for either the real or imaginary part of the matrix of ρ . This can be solved by defining the following alternative version of the \mathcal{CP} channel:

$$\mathcal{CP}'(\rho) = i \sum_{ijk=1}^3 \text{Tr}[\rho \text{SWAP} \sigma_i \sigma_j] \epsilon_{ijk} \sigma_k, \quad (\text{F18})$$

where we have multiplied by i to make the output matrix be Hermitian. One can readily check that this new CP channel is also equivariant; i.e.,

$$\mathcal{CP}'(U^{\otimes 2} \rho (U^\dagger)^{\otimes 2}) = U \mathcal{CP}'(\rho) U^\dagger. \quad (\text{F19})$$

Now, let us note that

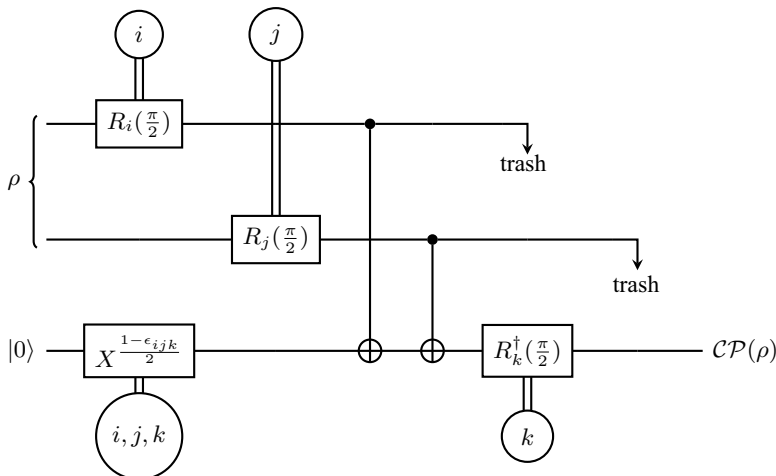
$$\text{SWAP}(YZ - ZY) = 2i(|\beta_{01}\rangle \langle \beta_{11}| + |\beta_{11}\rangle \langle \beta_{01}|), \quad (\text{F20})$$

$$\text{SWAP}(ZX - XZ) = 2(|\beta_{00}\rangle \langle \beta_{11}| - |\beta_{11}\rangle \langle \beta_{00}|), \quad (\text{F21})$$

$$\text{SWAP}(XY - YX) = -2i(|\beta_{10}\rangle \langle \beta_{11}| + |\beta_{11}\rangle \langle \beta_{10}|), \quad (\text{F22})$$

which means

$$\begin{aligned} \text{Tr}[\rho \text{SWAP}(YZ - ZY)] &= 4i\text{Re}[\rho_{24}], \\ \text{Tr}[\rho \text{SWAP}(ZX - XZ)] &= -4i\text{Im}[\rho_{14}], \\ \text{Tr}[\rho \text{SWAP}(XY - YX)] &= -4i\text{Re}[\rho_{34}]. \end{aligned} \quad (\text{F23})$$



Thus, we can now create combinations of the two versions of the \mathcal{CP} channels. For instance, the channel

$$\begin{aligned} \frac{1}{4}(\mathcal{CP}(\rho) + \mathcal{CP}'(\rho)) &= -(\text{Re}[a_{24}] - \text{Im}[a_{24}])X \\ &\quad + (\text{Re}[\rho_{14}] + \text{Im}[a_{14}])Y \\ &\quad + (\text{Re}[\rho_{34}] - \text{Im}[a_{34}])Z \end{aligned} \quad (\text{F24})$$

and

$$\begin{aligned} \frac{1}{4}(\mathcal{CP}(\rho) - \mathcal{CP}'(\rho)) &= (\text{Re}[a_{24}] + \text{Im}[a_{24}])X \\ &\quad + (\text{Re}[a_{14}] - \text{Im}[a_{14}])Y \\ &\quad - (\text{Re}[a_{34}] + \text{Im}[a_{34}])Z. \end{aligned} \quad (\text{F25})$$

By applying these channels (with the appropriate amount of completely depolarizing channel added such that the operation is physical), one may then recover the off-diagonal terms between the symmetric and antisymmetric subspaces.

Just as a curiosity, we note that the following combination (although not physical) is still interesting:

$$\frac{1}{4}(\mathcal{CP}(\rho) - i\mathcal{CP}'(\rho)) = i\rho_{24}^* X + \rho_{14}^* Y - i\rho_{34}^* Z. \quad (\text{F26})$$

In Fig. 14, we give a circuit for achieving the \mathcal{CP} channel in expectation. The intuition behind this circuit is that, without the Pauli rotations, this protocol changes the state of the ancilla such that $\text{Tr}[Z\rho_{\text{out}}] = \text{Tr}[ZZ\rho_{\text{in}}]$, where ρ_{out} is the state on the ancilla. Thus, by adding in Pauli rotations, we permute the Pauli group to embed $\text{Tr}[\sigma_i \sigma_j \rho]$ as $\text{Tr}[\sigma_k \rho_{\text{out}}]$ on the ancilla. Note that a similar circuit does not exist for \mathcal{CP}' , as $\text{SWAP} \otimes \sigma_i \sigma_j$ is not Hermitian. One could instead accomplish \mathcal{CP}' in expectation through the Hadamard test.

FIG. 14. The circuit for the cross-product channel. Here, $i, j, k \in \{X, Y, Z\}$. These classical random variables are drawn from $\{X, Y, Z\}$ uniformly and without replacement. That is, $\{i, j, k\} = \{X, Y, Z\}$ always but with the individual rotations randomly chosen. After performing this protocol, the reduced state on the ancilla will be equal, in expectation, to $\mathcal{CP}(\rho)$. Note that ϵ_{ijk} is the Levi-Civita symbol.

4. 1-to-2-qubit $SU(2)$ -equivariant maps

If we solve for the null space of the set of equivariant 1-to-2 maps, we see that all trace-preserving maps must have the following action. Say that our input is $\rho = \frac{1}{2}(\mathbb{1} + \boldsymbol{\sigma} \cdot \mathbf{r})$; then, the output state of the map will be

$$\begin{aligned} & \frac{1}{4} \mathbb{1} \otimes \mathbb{1} + \frac{a}{2} (XX + YY + ZZ) + \frac{b}{2} \mathbb{1} \otimes (\boldsymbol{\sigma} \cdot \mathbf{r}) \\ & + \frac{c}{2} (\boldsymbol{\sigma} \cdot \mathbf{r}) \otimes \mathbb{1} + \frac{d}{2} \mathbf{r} \cdot (YZ - ZY, ZX - XZ, XY - YX), \end{aligned} \tag{F27}$$

where a, b, c , and d are arbitrary real numbers (to be CP, there will be additional constraints). Note that this is effectively a linear combination of adjoints of the $2 \rightarrow 1$ maps, the first being Tr, the second SWAP, the third and fourth partial traces, and the final one the CP channel. This follows logically from representations of $SU(2)$ being self-dual. Recall that a linear map is equivariant if and only if its Choi operator satisfies $[(R^{\text{in}}(g))^* \otimes R^{\text{out}}(g), J^\phi] = 0$; i.e., J^ϕ lies in the commutant of the tensor product of the (dual) input and output representations. An equivariant $1 \mapsto 2$ map can then be associated with the commutant of

$$g^* \otimes g^{\otimes 2} \cong g^{\otimes 3}, \tag{F28}$$

which is exactly the same as the representation that commutes with the Choi operators of $2 \mapsto 1$ $SU(2)$ -equivariant channels.

[1] Y. LeCun, L. Bottou, Y. Bengio, and P. Haffner, Gradient-based learning applied to document recognition, *Proc. IEEE* **86**, 2278 (1998).

[2] M. M. Bronstein, J. Bruna, T. Cohen, and P. Veličković, Geometric deep learning: Grids, groups, graphs, geodesics, and gauges, *ArXiv:2104.13478*.

[3] E. J. Bekkers, M. W. Lafarge, M. Veta, K. A. Eppenhof, J. P. Pluim, and R. Duits, in *International Conference on Medical Image Computing and Computer-Assisted Intervention* (Springer, Granada, Spain, 2018), p. 440.

[4] S. Jegelka, Theory of graph neural networks: Representation and learning, *ArXiv:2204.07697*.

[5] R. Kondor, Z. Lin, and S. Trivedi, Clebsch-Gordan nets: A fully Fourier space spherical convolutional neural network, *Adv. Neural Inf. Process. Syst.* **31**, 10117 (2018).

[6] T. S. Cohen, M. Geiger, J. Köhler, and M. Welling, Spherical CNNs, *ArXiv:1801.10130*.

[7] A. Bogatskiy, B. Anderson, J. Offermann, M. Roussi, D. Miller, and R. Kondor, in *International Conference on Machine Learning* (PMLR, 2020), p. 992.

[8] M. Finzi, S. Stanton, P. Izmailov, and A. G. Wilson, in *International Conference on Machine Learning* (PMLR, 2020), p. 3165.

[9] B. Elesedy and S. Zaidi, in *International Conference on Machine Learning* (PMLR, 2021), p. 2959.

[10] T. Cohen and M. Welling, in *International Conference on Machine Learning* (PMLR, New York, New York, USA, 2016), p. 2990.

[11] R. Kondor and S. Trivedi, in *International Conference on Machine Learning* (PMLR, Stockholm, Stockholm Sweden, 2018), p. 2747.

[12] T. S. Cohen, M. Geiger, and M. Weiler, A general theory of equivariant CNNs on homogeneous spaces, *Adv. Neural Inf. Process. Syst.* **32**, 9142 (2019).

[13] T. Cohen, M. Weiler, B. Kicanaoglu, and M. Welling, in *International Conference on Machine Learning* (PMLR, Long Beach, California, USA, 2019), p. 1321.

[14] T. S. Cohen, Ph.D. thesis, Informatics Institute, Faculty of Science, University of Amsterdam, 2021, available online at <https://dare.uva.nl/search?identifier=0f7014ae-ee94-430e-a5d8-37d03d8d10e6>.

[15] J. Biamonte, P. Wittek, N. Pancotti, P. Rebentrost, N. Wiebe, and S. Lloyd, Quantum machine learning, *Nature* **549**, 195 (2017).

[16] M. Schuld, I. Sinayskiy, and F. Petruccione, The quest for a quantum neural network, *Quantum Inf. Process.* **13**, 2567 (2014).

[17] M. Schuld, I. Sinayskiy, and F. Petruccione, An introduction to quantum machine learning, *Contemp. Phys.* **56**, 172 (2015).

[18] M. Cerezo, A. Arrasmith, R. Babbush, S. C. Benjamin, S. Endo, K. Fujii, J. R. McClean, K. Mitarai, X. Yuan, L. Cincio, and P. J. Coles, Variational quantum algorithms, *Nat. Rev. Phys.* **3**, 625 (2021).

[19] M. Cerezo, G. Verdon, H.-Y. Huang, L. Cincio, and P. J. Coles, Challenges and opportunities in quantum machine learning, *Nat. Comput. Sci.* **2**, 567 (2022).

[20] J. Preskill, Quantum computing in the NISQ era and beyond, *Quantum* **2**, 79 (2018).

[21] H.-Y. Huang, M. Broughton, J. Cotler, S. Chen, J. Li, M. Mohseni, H. Neven, R. Babbush, R. Kueng, J. Preskill, and J. R. McClean, Quantum advantage in learning from experiments, *Science* **376**, 1182 (2022).

[22] Y. Liu, S. Arunachalam, and K. Temme, A rigorous and robust quantum speed-up in supervised machine learning, *Nat. Phys.* **2**, 567 (2022).

[23] I. Cong, S. Choi, and M. D. Lukin, Quantum convolutional neural networks, *Nat. Phys.* **15**, 1273 (2019).

[24] L. Schatzki, A. Arrasmith, P. J. Coles, and M. Cerezo, Entangled datasets for quantum machine learning, *ArXiv:2109.03400*.

[25] M. C. Caro, H.-Y. Huang, M. Cerezo, K. Sharma, A. Sornborger, L. Cincio, and P. J. Coles, Generalization in quantum machine learning from few training data, *Nat. Commun.* **13**, 4919 (2022).

[26] J. Kübler, S. Buchholz, and B. Schölkopf, The inductive bias of quantum kernels, *Adv. Neural Inf. Process. Syst.* **34**, 12661 (2021).

[27] J. R. McClean, S. Boixo, V. N. Smelyanskiy, R. Babbush, and H. Neven, Barren plateaus in quantum neural network training landscapes, *Nat. Commun.* **9**, 1 (2018).

[28] M. Cerezo, A. Sone, T. Volkoff, L. Cincio, and P. J. Coles, Cost function dependent barren plateaus in shallow parametrized quantum circuits, *Nat. Commun.* **12**, 1 (2021).

- [29] K. Sharma, M. Cerezo, L. Cincio, and P. J. Coles, Trainability of dissipative perceptron-based quantum neural networks, *Phys. Rev. Lett.* **128**, 180505 (2022).
- [30] Z. Holmes, A. Arrasmith, B. Yan, P. J. Coles, A. Albrecht, and A. T. Sornborger, Barren plateaus preclude learning scramblers, *Phys. Rev. Lett.* **126**, 190501 (2021).
- [31] Z. Holmes, K. Sharma, M. Cerezo, and P. J. Coles, Connecting ansatz expressibility to gradient magnitudes and barren plateaus, *PRX Quantum* **3**, 010313 (2022).
- [32] M. Cerezo and P. J. Coles, Higher order derivatives of quantum neural networks with barren plateaus, *Quantum Sci. Technol.* **6**, 035006 (2021).
- [33] A. Arrasmith, M. Cerezo, P. Czarnik, L. Cincio, and P. J. Coles, Effect of barren plateaus on gradient-free optimization, *Quantum* **5**, 558 (2021).
- [34] A. Arrasmith, Z. Holmes, M. Cerezo, and P. J. Coles, Equivalence of quantum barren plateaus to cost concentration and narrow gorges, *Quantum Sci. Technol.* **7**, 045015 (2022).
- [35] C. O. Marrero, M. Kieferová, and N. Wiebe, Entanglement-induced barren plateaus, *PRX Quantum* **2**, 040316 (2021).
- [36] T. L. Patti, K. Najafi, X. Gao, and S. F. Yelin, Entanglement devised barren plateau mitigation, *Phys. Rev. Res.* **3**, 033090 (2021).
- [37] A. Uvarov and J. D. Biamonte, On barren plateaus and cost function locality in variational quantum algorithms, *J. Phys. A: Math. Theor.* **54**, 245301 (2021).
- [38] S. Thanasilp, S. Wang, N. A. Nghiem, P. J. Coles, and M. Cerezo, Subtleties in the trainability of quantum machine learning models, *ArXiv:2110.14753*.
- [39] M. Larocca, P. Czarnik, K. Sharma, G. Muraleedharan, P. J. Coles, and M. Cerezo, Diagnosing barren plateaus with tools from quantum optimal control, *Quantum* **6**, 824 (2022).
- [40] S. Wang, E. Fontana, M. Cerezo, K. Sharma, A. Sone, L. Cincio, and P. J. Coles, Noise-induced barren plateaus in variational quantum algorithms, *Nat. Commun.* **12**, 1 (2021).
- [41] R. Wiersema, C. Zhou, Y. de Sereville, J. F. Carrasquilla, Y. B. Kim, and H. Yuen, Exploring entanglement and optimization within the Hamiltonian variational ansatz, *PRX Quantum* **1**, 020319 (2020).
- [42] R. Horodecki, P. Horodecki, M. Horodecki, and K. Horodecki, Quantum entanglement, *Rev. Mod. Phys.* **81**, 865 (2009).
- [43] S. Szalay, Multipartite entanglement measures, *Phys. Rev. A* **92**, 042329 (2015).
- [44] M. Larocca, F. Sauvage, F. M. Sbahi, G. Verdon, P. J. Coles, and M. Cerezo, Group-invariant quantum machine learning, *PRX Quantum* **3**, 030341 (2022).
- [45] J. J. Meyer, M. Mularski, E. Gil-Fuster, A. A. Mele, F. Arzani, A. Wilms, and J. Eisert, Exploiting symmetry in variational quantum machine learning, *PRX Quantum* **4**, 010328 (2023).
- [46] P. Memyei, K. Meichanetzidis, and I. I. Ceylan, Equivariant quantum graph circuits, *ArXiv:2112.05261*.
- [47] A. Skolik, M. Cattelan, S. Yarkoni, T. Bäck, and V. Dunjko, Equivariant quantum circuits for learning on weighted graphs, *npj Quantum Inf.* **9**, 47 (2023).
- [48] H. Zheng, Z. Li, J. Liu, S. Strelchuk, and R. Kondor, Speeding up learning quantum states through group equivariant convolutional quantum ansätze, *PRX Quantum* **4**, 020327 (2023).
- [49] H. Zheng, Z. Li, J. Liu, S. Strelchuk, and R. Kondor, On the super-exponential quantum speedup of equivariant quantum machine learning algorithms with $SU(d)$ symmetry, *ArXiv:2207.07250*.
- [50] H. Zheng, G. S. Ravi, H. Wang, K. Setia, F. T. Chong, and J. Liu, Benchmarking variational quantum circuits with permutation symmetry, *ArXiv:2211.12711*.
- [51] Z. Li, H. Zheng, E. Thiede, J. Liu, and R. Kondor, Group-equivariant neural networks with fusion diagrams, *ArXiv:2211.07482*.
- [52] G. Verdon, T. McCourt, E. Luzhnica, V. Singh, S. Leichenauer, and J. Hidary, Quantum graph neural networks, *ArXiv:1909.12264*.
- [53] B. Anderson, T. S. Hy, and R. Kondor, Cormorant: Covariant molecular neural networks, *Adv. Neural Inf. Process. Syst.* **32**, 14537 (2019).
- [54] M. Geiger and T. Smidt, E3NN: Euclidean neural networks, *ArXiv:2207.09453*.
- [55] N. Thomas, T. Smidt, S. Kearnes, L. Yang, L. Li, K. Kohlhoff, and P. Riley, Tensor field networks: Rotation- and translation-equivariant neural networks for 3D point clouds, *ArXiv:1802.08219* (2018).
- [56] Z. Chen, N. Andrejevic, T. Smidt, Z. Ding, Q. Xu, Y.-T. Chi, Q. T. Nguyen, A. Alatas, J. Kong, and M. Li, Direct prediction of phonon density of states with euclidean neural networks, *Adv. Sci.* **8**, 2004214 (2021).
- [57] H. Pan and R. Kondor, in *Proceedings of The 25th International Conference on Artificial Intelligence and Statistics*, edited by G. Camps-Valls, F. J. R. Ruiz, and I. Valera, Proceedings of Machine Learning Research, Vol. 151 (PMLR, 2022), p. 5987.
- [58] E. H. Thiede, T. S. Hy, and R. Kondor, The general theory of permutation equivariant neural networks and higher order graph variational encoders, *ArXiv:2004.03990*.
- [59] T. S. Cohen and M. Welling, Steerable CNNs, *ArXiv:1612.08498*.
- [60] M. Finzi, M. Welling, and A. G. Wilson, in *International Conference on Machine Learning* (PMLR, 2021), p. 3318.
- [61] J. Brandstetter, R. van den Berg, M. Welling, and J. K. Gupta, Clifford neural layers for PDE modeling, *ArXiv:2209.04934*.
- [62] Z. Li, Y. Zhang, and S. Arora, Why are convolutional nets more sample-efficient than fully-connected nets? *ArXiv:2010.08515*.
- [63] A. Sannai, M. Imaizumi, and M. Kawano, in *Uncertainty in Artificial Intelligence* (PMLR, 2021), p. 771.
- [64] J. Sokolic, R. Giryas, G. Sapiro, and M. Rodrigues, in *Artificial Intelligence and Statistics* (PMLR, Ft. Lauderdale, FL, USA, 2017), p. 1094.
- [65] H. Maron, E. Fetaya, N. Segol, and Y. Lipman, in *Proceedings of the 36th International Conference on Machine Learning*, edited by K. Chaudhuri and R. Salakhutdinov, Proceedings of Machine Learning Research, Vol. 97 (PMLR, Long Beach, California, USA, 2019), p. 4363.
- [66] D. Yarotsky, Universal approximations of invariant maps by neural networks, *Constr. Approx.* **55**, 407 (2022).

- [67] N. Keriven and G. Peyré, in *Advances in Neural Information Processing Systems*, edited by H. Wallach, H. Larochelle, A. Beygelzimer, F. d' Alché-Buc, E. Fox, and R. Garnett (Curran Associates, Inc., Vancouver Canada, 2019), Vol. 32.
- [68] S. Ravanbakhsh, in *International Conference on Machine Learning* (PMLR, 2020), p. 7996.
- [69] S. Villar, D. W. Hogg, K. Storey-Fisher, W. Yao, and B. Blum-Smith, Scalars are universal: Equivariant machine learning, structured like classical physics, *Adv. Neural Inf. Process. Syst.* **34**, 28848 (2021).
- [70] The alternative term *covariance* is used in some quantum information theory literature. Here, we will adopt the terminology *equivariance* for cohesiveness with the geometric deep-learning literature.
- [71] A. Holevo, Additivity conjecture and covariant channels, *Int. J. Quantum Inf.* **3**, 41 (2005).
- [72] W. G. Ritter, Quantum channels and representation theory, *J. Math. Phys.* **46**, 082103 (2005).
- [73] M. B. Hastings, Superadditivity of communication capacity using entangled inputs, *Nat. Phys.* **5**, 255 (2009).
- [74] M. M. Wilde, M. Tomamichel, and M. Berta, Converse bounds for private communication over quantum channels, *IEEE Trans. Inf. Theory* **63**, 1792 (2017).
- [75] R. Koenig and S. Wehner, A strong converse for classical channel coding using entangled inputs, *Phys. Rev. Lett.* **103**, 070504 (2009).
- [76] N. Datta, M. Tomamichel, and M. M. Wilde, On the second-order asymptotics for entanglement-assisted communication, *Quantum Inf. Process.* **15**, 2569 (2016).
- [77] F. Leditzky, E. Kaur, N. Datta, and M. M. Wilde, Approaches for approximate additivity of the Holevo information of quantum channels, *Phys. Rev. A* **97**, 012332 (2018).
- [78] M. Gschwendtner, A. Bluhm, and A. Winter, Programmability of covariant quantum channels, *Quantum* **5**, 488 (2021).
- [79] D. Grinko and M. Ozols, Linear programming with unitary-equivariant constraints, [ArXiv:2207.05713](https://arxiv.org/abs/2207.05713).
- [80] L. Kong and Z.-W. Liu, Near-optimal covariant quantum error-correcting codes from random unitaries with symmetries, *PRX Quantum* **3**, 020314 (2022).
- [81] P. Faist, S. Nezami, V. V. Albert, G. Salton, F. Pastawski, P. Hayden, and J. Preskill, Continuous symmetries and approximate quantum error correction, *Phys. Rev. X* **10**, 041018 (2020).
- [82] S. Zhou, Z.-W. Liu, and L. Jiang, New perspectives on covariant quantum error correction, *Quantum* **5**, 521 (2021).
- [83] G. M. D'ariano, Extremal covariant quantum operations and positive operator valued measures, *J. Math. Phys.* **45**, 3620 (2004).
- [84] M. Hayashi, *Group Representation for Quantum Theory* (Springer, Switzerland, 2017).
- [85] H. Krovi, An efficient high dimensional quantum Schur transform, *Quantum* **3**, 122 (2019).
- [86] A. W. Harrow, Applications of coherent classical communication and the Schur transform to quantum information theory, [ArXiv:quant-ph/0512255](https://arxiv.org/abs/quant-ph/0512255).
- [87] D. Bacon, I. L. Chuang, and A. W. Harrow, Efficient quantum circuits for Schur and Clebsch-Gordan transforms, *Phys. Rev. Lett.* **97**, 170502 (2006).
- [88] I. Marvian, Restrictions on realizable unitary operations imposed by symmetry and locality, *Nat. Phys.* **18**, 283 (2022).
- [89] I. Marvian, H. Liu, and A. Hulse, Rotationally-invariant circuits: Universality with the exchange interaction and two ancilla qubits, [ArXiv:2202.01963](https://arxiv.org/abs/2202.01963).
- [90] A. Hulse, H. Liu, and I. Marvian, Qudit circuits with $SU(d)$ symmetry: Locality imposes additional conservation laws, [ArXiv:2105.12877](https://arxiv.org/abs/2105.12877).
- [91] A. M. Childs and W. van Dam, Quantum algorithms for algebraic problems, *Rev. Mod. Phys.* **82**, 1 (2010).
- [92] M. Al Nuwairan, The extreme points of $SU(2)$ -irreducibly covariant channels, *Int. J. Math.* **25**, 1450048 (2014).
- [93] M. Mozrzyk, M. Studziński, and N. Datta, Structure of irreducibly covariant quantum channels for finite groups, *J. Math. Phys.* **58**, 052204 (2017).
- [94] L. Memarzadeh and B. C. Sanders, Group-covariant extreme and quasiextreme channels, *Phys. Rev. Res.* **4**, 033206 (2022).
- [95] S. Kazi, M. Larocca, and M. Cerezo, On the universality of s_n -equivariant k -body gates, [ArXiv:2303.00728](https://arxiv.org/abs/2303.00728).
- [96] L. Schatzki, M. Larocca, Q. T. Nguyen, F. Sauvage, and M. Cerezo, Theoretical guarantees for permutation-equivariant quantum neural networks, *npj Quantum Inf.* **10**, 12 (2024).
- [97] F. Sauvage, M. Larocca, P. J. Coles, and M. Cerezo, Building spatial symmetries into parameterized quantum circuits for faster training, [ArXiv:2207.14413](https://arxiv.org/abs/2207.14413).
- [98] L. Cincio, Y. Subaşı, A. T. Sornborger, and P. J. Coles, Learning the quantum algorithm for state overlap, *New J. Phys.* **20**, 113022 (2018).
- [99] S. Van Enk and C. Beenakker, Measuring $\text{Tr } \rho^n$ on single copies of ρ using random measurements, *Phys. Rev. Lett.* **108**, 110503 (2012).
- [100] H.-Y. Huang, R. Kueng, and J. Preskill, Predicting many properties of a quantum system from very few measurements, *Nat. Phys.* **16**, 1050 (2020).
- [101] B. Simon, *Representations of Finite and Compact Groups*, Graduate Studies in Mathematics, Vol. 10 (American Mathematical Society, 1996).
- [102] W. Fulton and J. Harris, *Representation Theory: A First Course* (Springer, Springer New York, NY, 1991).
- [103] R. Goodman and N. R. Wallach, *Representations and Invariants of the Classical Groups* (Cambridge University Press, Cambridge, UK, 2000).
- [104] M. Ragone, Q. T. Nguyen, L. Schatzki, P. Braccia, M. Larocca, F. Sauvage, P. J. Coles, and M. Cerezo, Representation theory for geometric quantum machine learning, [ArXiv:2210.07980](https://arxiv.org/abs/2210.07980).
- [105] J. S. Otterbach, *et al.*, Unsupervised machine learning on a hybrid quantum computer, [ArXiv:1712.05771](https://arxiv.org/abs/1712.05771).
- [106] I. Kerenidis, J. Landman, A. Luongo, and A. Prakash, q-means: A quantum algorithm for unsupervised machine learning, *Adv. Neural Inf. Process. Syst.* **32**, 4134 (2019).
- [107] P.-L. Dallaire-Demers and N. Killoran, Quantum generative adversarial networks, *Phys. Rev. A* **98**, 012324 (2018).

- [108] M. Benedetti, D. Garcia-Pintos, O. Perdomo, V. Leyton-Ortega, Y. Nam, and A. Perdomo-Ortiz, A generative modeling approach for benchmarking and training shallow quantum circuits, *npj Quantum Inf.* **5**, 45 (2019).
- [109] M. Kieferova, O. M. Carlos, and N. Wiebe, Quantum generative training using Rényi divergences, *ArXiv:2106.09567*.
- [110] J. Romero and A. Aspuru-Guzik, Variational quantum generators: Generative adversarial quantum machine learning for continuous distributions, *Adv. Quantum Technol.* **4**, 2000003 (2021).
- [111] V. Saggio, B. E. Asenbeck, A. Hamann, T. Strömberg, P. Schiаны, V. Dunjko, N. Friis, N. C. Harris, M. Hochberg, D. Englund, S. Wölk, H. J. Briegel, and P. Walther, Experimental quantum speed-up in reinforcement learning agents, *Nature* **591**, 229 (2021).
- [112] A. Skolik, S. Jerbi, and V. Dunjko, Quantum agents in the gym: A variational quantum algorithm for deep q -learning, *ArXiv:2103.15084*.
- [113] V. Havlíček, A. D. Córcoles, K. Temme, A. W. Harrow, A. Kandala, J. M. Chow, and J. M. Gambetta, Supervised learning with quantum-enhanced feature spaces, *Nature* **567**, 209 (2019).
- [114] B. C. Hall, *Lie Groups, Lie Algebras, and Representations* (Springer, Springer New York, NY, 2013).
- [115] G. B. Folland, *A Course in Abstract Harmonic Analysis* (CRC Press, Boca Raton, Florida, USA, 2016), Vol. 29.
- [116] M. M. Wilde, *Quantum Information Theory* (Cambridge University Press, Cambridge, UK, 2013).
- [117] J. Watrous, *The Theory of Quantum Information* (Cambridge University Press, Cambridge, UK, 2018).
- [118] S. Sim, P. D. Johnson, and A. Aspuru-Guzik, Expressibility and entangling capability of parameterized quantum circuits for hybrid quantum-classical algorithms, *Adv. Quantum Technol.* **2**, 1900070 (2019).
- [119] G. Castelazo, Q. T. Nguyen, G. De Palma, D. Englund, S. Lloyd, and B. T. Kiani, Quantum algorithms for group convolution, cross-correlation, and equivariant transformations, *ArXiv:2109.11330*.
- [120] A. Gilyén, Y. Su, G. H. Low, and N. Wiebe, Quantum singular value transformation and beyond: exponential improvements for quantum matrix arithmetics, in *Proceedings of the 51st Annual ACM SIGACT Symposium on Theory of Computing* (ACM Publications, Phoenix AZ USA, 2019), p. 193.
- [121] Z. Holmes, N. Coble, A. T. Sornborger, and Y. Subaşı, On nonlinear transformations in quantum computation, *ArXiv:2112.12307*.
- [122] J. R. Glick, T. P. Gujarati, A. D. Córcoles, Y. Kim, A. Kandala, J. M. Gambetta, and K. Temme, Covariant quantum kernels for data with group structure, *ArXiv:2105.03406*.
- [123] R. Zeier and T. Schulte-Herbrüggen, Symmetry principles in quantum systems theory, *J. Math. Phys.* **52**, 113510 (2011).
- [124] C. J. Wood, J. D. Biamonte, and D. G. Cory, Tensor networks and graphical calculus for open quantum systems, *ArXiv:1111.6950*.
- [125] R. A. Horn and C. R. Johnson, *Topics in Matrix Analysis* (Cambridge University Press, Cambridge, UK, 1991).
- [126] L. N. Trefethen and D. Bau III, *Numerical Linear Algebra* (SIAM, Philadelphia, PA, USA, 1997), Vol. 50.
- [127] B. Collins and P. Śniady, Integration with respect to the Haar measure on unitary, orthogonal and symplectic group, *Commun. Math. Phys.* **264**, 773 (2006).
- [128] Z. Puchala and J. A. Miszczak, Symbolic integration with respect to the Haar measure on the unitary groups, *Bull. Pol. Acad. Sci. Tech. Sci.* **65**, 21 (2017).
- [129] G. Tóth and J. J. Garcia-Ripoll, Efficient algorithm for multiqubit twirling for ensemble quantum computation, *Phys. Rev. A* **75**, 042311 (2007).
- [130] The set of equivariant CPTP channels forms a convex hull. Channels on the boundary of this hull are said to be extremal and fully characterize the convex hull. However, this set of extremal channels may not be a finite set.
- [131] W. Bosma, J. Cannon, and C. Playoust, The Magma algebra system. I. The user language, *J. Symbolic Comput.* **24**, 235 (1997), *computational algebra and number theory (London, 1993)*.
- [132] M. A. Van Leeuwen, A. M. Cohen, and B. Lisser, LIE: A package for lie group computations (Centrum voor Wiskunde en Informatica, Amsterdam, 1992), p. 113.
- [133] K. Bharti and T. Haug, Quantum-assisted simulator, *Phys. Rev. A* **104**, 042418 (2021).
- [134] F. T. Chong, D. Franklin, and M. Martonosi, Programming languages and compiler design for realistic quantum hardware, *Nature* **549**, 180 (2017).
- [135] T. Häner, D. S. Steiger, K. Svore, and M. Troyer, A software methodology for compiling quantum programs, *Quantum Sci. Technol.* **3**, 020501 (2018).
- [136] K. Sharma, S. Khatri, M. Cerezo, and P. J. Coles, Noise resilience of variational quantum compiling, *New J. Phys.* **22**, 043006 (2020).
- [137] I. MacCormack, C. Delaney, A. Galda, N. Aggarwal, and P. Narang, Branching quantum convolutional neural networks, *ArXiv:2012.14439*.
- [138] L. Franken and B. Georgiev, in *28th European Symposium on Artificial Neural Networks, Computational Intelligence and Machine Learning* (2020).
- [139] A. Pesah, M. Cerezo, S. Wang, T. Volkoff, A. T. Sornborger, and P. J. Coles, Absence of barren plateaus in quantum convolutional neural networks, *Phys. Rev. X* **11**, 041011 (2021).
- [140] J. B. Parkinson and D. J. Farnell, *An Introduction to Quantum Spin Systems* (Springer, Springer Berlin, Heidelberg, 2010), Vol. 816.
- [141] A. Elben, J. Yu, G. Zhu, M. Hafezi, F. Pollmann, P. Zoller, and B. Vermersch, Many-body topological invariants from randomized measurements in synthetic quantum matter, *Sci. Adv.* **6**, eaaz3666 (2020).
- [142] H.-Y. Huang, R. Kueng, G. Torlai, V. V. Albert, and J. Preskill, Provably efficient machine learning for quantum many-body problems, *Science* **377**, eabk3333 (2022).
- [143] J. L. Beckey, N. Gigena, P. J. Coles, and M. Cerezo, Computable and operationally meaningful multipartite entanglement measures, *Phys. Rev. Lett.* **127**, 140501 (2021).
- [144] L. Schatzki, G. Liu, M. Cerezo, and E. Chitambar, A hierarchy of multipartite correlations based on concentratable entanglement, *ArXiv:2209.07607*.
- [145] R. Goodman and N. R. Wallach, *Symmetry, Representations, and Invariants* (Springer, Springer New York, NY, 2009), Vol. 255.

- [146] F. D. M. Haldane, Nonlinear field theory of large-spin Heisenberg antiferromagnets: Semiclassically quantized solitons of the one-dimensional easy-axis Néel state, *Phys. Rev. Lett.* **50**, 1153 (1983).
- [147] S. Sachdev, Quantum phase transitions, *Handbook of Magnetism and Advanced Magnetic Materials* (2007).
- [148] M. Cerezo, R. Rossignoli, N. Canosa, and E. Ríos, Factorization and criticality in finite XXZ systems of arbitrary spin, *Phys. Rev. Lett.* **119**, 220605 (2017).
- [149] T. Volkoff and P. J. Coles, Large gradients via correlation in random parameterized quantum circuits, *Quantum Sci. Technol.* **6**, 025008 (2021).
- [150] N. Srivastava, G. Hinton, A. Krizhevsky, I. Sutskever, and R. Salakhutdinov, DROPOUT: A simple way to prevent neural networks from overfitting, *J. Mach. Learning Res.* **15**, 1929 (2014).
- [151] F. Pollmann and A. M. Turner, Detection of symmetry-protected topological phases in one dimension, *Phys. Rev. B* **86**, 125441 (2012).
- [152] A. Kandala, A. Mezzacapo, K. Temme, M. Takita, M. Brink, J. M. Chow, and J. M. Gambetta, Hardware-efficient variational quantum eigensolver for small molecules and quantum magnets, *Nature* **549**, 242 (2017).
- [153] D. P. Kingma and J. Ba, Adam: A method for stochastic optimization, *ArXiv:1412.6980* (2014).
- [154] L. Bittel and M. Kliesch, Training variational quantum algorithms is NP-hard, *Phys. Rev. Lett.* **127**, 120502 (2021).
- [155] E. R. Anschuetz and B. T. Kiani, Beyond barren plateaus: Quantum variational algorithms are swamped with traps, *Nat. Commun.* **13**, 7760 (2022).
- [156] E. Fontana, M. Cerezo, A. Arrasmith, I. Rungger, and P. J. Coles, Non-trivial symmetries in quantum landscapes and their resilience to quantum noise, *Quantum* **6**, 804 (2022).
- [157] E. R. Anschuetz, Critical Points in Quantum Generative Models, in *International Conference on Learning Representations* (2022).
- [158] M. Larocca, N. Ju, D. García-Martín, P. J. Coles, and M. Cerezo, Theory of overparametrization in quantum neural networks, *Nat. Comput. Sci.* **3**, 542 (2023).
- [159] H.-Y. Huang, M. Broughton, M. Mohseni, R. Babbush, S. Boixo, H. Neven, and J. R. McClean, Power of data in quantum machine learning, *Nat. Commun.* **12**, 1 (2021).
- [160] L. Banchi, J. Pereira, and S. Pirandola, Generalization in quantum machine learning: A quantum information standpoint, *PRX Quantum* **2**, 040321 (2021).
- [161] K. Bu, D. E. Koh, L. Li, Q. Luo, and Y. Zhang, Effects of quantum resources on the statistical complexity of quantum circuits, *ArXiv:2102.03282*.
- [162] S. Thanasilp, S. Wang, M. Cerezo, and Z. Holmes, Exponential concentration and untrainability in quantum kernel methods, *ArXiv:2208.11060*.
- [163] R. Wang, R. Walters, and R. Yu, Approximately equivariant networks for imperfectly symmetric dynamics, *ArXiv:2201.11969*.
- [164] I. N. M. Le, O. Kiss, J. Schuhmacher, I. Tavernelli, and F. Tacchino, Symmetry-invariant quantum machine learning force fields, *ArXiv:2311.11362*.
- [165] M. T. West, M. Seviar, and M. Usman, Reflection equivariant quantum neural networks for enhanced image classification, *Mach. Learning: Sci. Technol.* **4**, 035027 (2023).
- [166] S. Y. Chang, M. Grossi, B. Le Saux, and S. Vallecorsa, in *2023 IEEE International Conference on Quantum Computing and Engineering (QCE)* (Bellevue, Washington, USA, 2023), Vol. 1, p. 229.
- [167] M. T. West, J. Heredge, M. Seviar, and M. Usman, Provably trainable rotationally equivariant quantum machine learning, *ArXiv:2311.05873*.
- [168] R. D. East, G. Alonso-Linaje, and C.-Y. Park, All you need is spin: $SU(2)$ equivariant variational quantum circuits based on spin networks, *ArXiv:2309.07250*.
- [169] S. Das, S. Martina, and F. Caruso, The role of data embedding in equivariant quantum convolutional neural networks, *ArXiv:2312.13250*.
- [170] C. Tüysüz, S. Y. Chang, M. Demidik, K. Jansen, S. Vallecorsa, and M. Grossi, Symmetry breaking in geometric quantum machine learning in the presence of noise, *ArXiv:2401.10293*.
- [171] Z. Webb, The Clifford group forms a unitary 3-design, *Quantum Inf. Comput.* **16**, 1379 (2016).
- [172] B. Collins, S. Matsumoto, and J. Novak, The Weingarten calculus, *Not. Am. Math. Soc.* **69**, 734 (2022).
- [173] K. Koike, On the decomposition of tensor products of the representations of the classical groups: By means of the universal characters, *Adv. Math. (NY)* **74**, 57 (1989).
- [174] R. Iten, O. Reardon-Smith, E. Malvetti, L. Mondada, G. Pauvert, E. Redmond, R. S. Kohli, and R. Colbeck, Introduction to UNIVERSALQCOMPILER, *ArXiv:1904.01072* (2019).
- [175] R. Iten, R. Colbeck, I. Kukuljan, J. Home, and M. Christandl, Quantum circuits for isometries, *Phys. Rev. A* **93**, 032318 (2016).
- [176] R. Iten, R. Colbeck, and M. Christandl, Quantum circuits for quantum channels, *Phys. Rev. A* **95**, 052316 (2017).
- [177] S. Khatri, R. LaRose, A. Poremba, L. Cincio, A. T. Sornborger, and P. J. Coles, Quantum-assisted quantum compiling, *Quantum* **3**, 140 (2019).
- [178] L. Cincio, K. Rudinger, M. Sarovar, and P. J. Coles, Machine learning of noise-resilient quantum circuits, *PRX Quantum* **2**, 010324 (2021).