

Permutation-equivariant quantum convolutional neural networks

Sreetama Das^{1,3}, Filippo Caruso^{1,2,3}

¹*Istituto Nazionale di Ottica del Consiglio Nazionale delle Ricerche (CNR-INO), I-50019 Sesto Fiorentino, Italy*

²*Department of Physics and Astronomy, University of Florence, Via Sansone 1, Sesto Fiorentino, I-50019, Italy*

³*European Laboratory for Non-Linear Spectroscopy (LENs),
University of Florence, Via Nello Carrara 1, Sesto Fiorentino, I-50019, Italy*

The Symmetric group S_n manifests itself in large classes of quantum systems as the invariance of certain characteristics of a quantum state with respect to permuting the qubits. The subgroups of S_n arise, among many other contexts, to describe label symmetry of classical images with respect to spatial transformations, e.g. reflection or rotation. Equipped with the formalism of geometric quantum machine learning, in this work we propose the architectures of equivariant quantum convolutional neural networks (EQCNNs) adherent to S_n and its subgroups. We demonstrate that a careful choice of pixel-to-qubit embedding order can facilitate easy construction of EQCNNs for small subgroups of S_n . Our novel EQCNN architecture corresponding to the full permutation group S_n is built by applying all possible QCNNs with equal probability, which can also be conceptualized as a dropout strategy in quantum neural networks. For subgroups of S_n , our numerical results using MNIST datasets show better classification accuracy than non-equivariant QCNNs. The S_n -equivariant QCNN architecture shows significantly improved training and test performance than non-equivariant QCNN for classification of connected and non-connected graphs. When trained with sufficiently large number of data, the S_n -equivariant QCNN shows better average performance compared to S_n -equivariant QNN. These results contribute towards building powerful quantum machine learning architectures in permutation-symmetric systems.

I. INTRODUCTION

The field at the interface of quantum technology and machine learning has been subject to extensive research in recent times. A significant fraction of these are aimed at building machine learning models using quantum systems. Of particular interest are quantum neural networks (QNNs) [1–5], analogous to classical neural networks in deep learning. The central component in a prototypical QNN is a quantum circuit with single and multiple qubit gates with trainable parameters. It has been widely known as parametric quantum circuit (PQC) or variational quantum circuit (VQC) [6–10]. PQC-based QNNs have been used to design quantum analogs of well-known classical machine learning architectures, e.g. quantum autoencoder [8], quantum convolutional neural networks (QCNNs) [11–13], quantum generative adversarial networks (QGANs) [14–19], quantum generative diffusion models [20–22] etc. The cost function obtained from the above networks is optimized using a classical optimization routine, thus they are *hybrid quantum-classical networks*. They are designed to function with both classical and quantum data. Despite the potential speed-up and signatures of success of quantum machine learning over classical ML models [23–27], there exist several limitations for QNNs, the *barren plateau* problem being one of them [28–31]. On one hand, an arbitrary *hardware-efficient* PQC ansatz is expected to have high expressibility [32] to ensure that the solution of the optimization problem is close enough to the actual solution. On the other hand, PQCs with higher expressibility are increasingly prone to exhibiting barren plateau [28, 29]. It is therefore a crucial task to mitigate barren plateau for a practical application of quantum neural networks.

One way to improve the trainability and generalization in machine learning algorithms is to introduce *inductive bias* in the learning model, i.e. to use some prior known information about the dataset to build a problem-specific model and

constrain the optimization space of the network. Particularly, *geometric machine learning* refers to a scheme in which the known symmetries of the dataset are used to construct a network which respects those symmetries [33]. In past few years, the theory of geometric quantum machine learning has been developed by a number of works [34–40]. The performance of symmetry-enhanced QNNs, also called Equivariant QNNs (EQNNs), has been investigated for a number of commonly encountered symmetries in classical datasets as well as quantum systems. These studies show that some classes of EQNNs are devoid of the barren plateau problem [41–43]. Moreover, the performance of an EQNN shows improvement over a non-equivariant QNN for pattern recognition and image classification [37, 43–46].

Among all the symmetry groups, the group of all permutations of n objects, which is also called Symmetric group (S_n), is a crucial one. It is known that every group is isomorphic to some subgroup of a Symmetric group. Symmetric group frequently arises in physical scenarios where certain properties of a quantum state remain invariant under arbitrary permutation of the qubits. To exemplify, the genuine multi-party entangled quantum states remain genuinely entangled when the subsystems are permuted, or the set of connected graphs remain connected under permutation of the vertices. There also exist scenarios for which the permutation symmetry in a system corresponds to a subgroup of S_n . For example, the states of a many-body quantum system in presence of periodic boundary condition remain unchanged under cyclic permutation of qubits, which corresponds to cyclic group Z_n . Another example is the label symmetry of classical images, i.e. the class labels of images remain unchanged under spatial transformations like reflection or rotation, which corresponds to permuting the pixels in a certain way. The S_n -equivariant QNN has been studied in [42], with a theoretical proof of absence of barren plateau. For small subgroups of S_n relevant to classical images, EQNN can be easily built by visualizing

the qubits as pixels in a 2D plane [37].

While QNN broadly refers to a general PQC architecture, in this work, we focus our attention to a specific type of QNN, namely quantum convolutional neural networks. We propose an architecture of equivariant quantum convolutional neural network (EQCNN) for Symmetric group and its subgroups, which we jointly refer to as *permutation group*. Our architecture is based on the well-known model of QCNN introduced in Ref. [11], which shows a considerably good performance for classifying topological phases of many-body quantum systems [11, 47–50] as well as classical images [51]. A few works have studied EQCNN for classical images [45, 46], all of them using the quantum amplitude embedding. For reflection or rotation symmetry of images, amplitude embedding results in a representation of permutation group that acts locally on each qubit. However, when using qubit embedding (also known as angle embedding), the resulting representations correspond to subgroups of S_n and act non-locally on multiple qubits, as we will elaborate in the next section. In addition to that, the QCNN proposed in Ref. [11] is subject to some structural restrictions, i.e. the translational invariance of the convolution layers and the qubit-reducing feature of pooling layers. The main challenge in designing EQCNN is to comply with these constraints under non-local action of the unitary representations in general. In this work, we try to resolve this issue. We demonstrate that for subgroups of S_n , it is important to choose a particular order of pixel-to-qubit mapping which enables one to build the EQCNN ansatz in a translationally-invariant manner. Equipped with this, we put forward EQCNN architectures for reflection symmetry and $\pi/2$ -rotational symmetry of images, which shows improved performance compared to non-equivariant QCNNs for classifying MNIST datasets. For full permutation symmetry of S_n , our EQCNN is built by applying all QCNNs at random with equal probability. Such a probabilistic picture of QNN, although mentioned briefly in Ref. [38], has not been explored significantly before. We build a quantum circuit to realize this stochastic action of the convolutional and pooling layers. We test its performance for classification of connected and non-connected graphs against non-equivariant QCNN as well as S_n -equivariant QNN. While the EQCNN always performs better than the former, for the latter it shows an improvement when trained with comparatively large number of data-points.

II. METHODS

A. Quantum Convolutional Neural Network (QCNN)

The QCNN architecture of ref. [11] is inspired from the multi-scale entanglement renormalization ansatz (MERA) representation of quantum many-body states [52]. Any QNN has three components: an n -qubit quantum register encoding the input quantum state $|\Psi\rangle$, a PQC $\mathcal{U}_\theta$ acting on $|\Psi\rangle$, and lastly a measurement M on some (or all) of the qubits. In case of QCNN, $\mathcal{U}_\theta$ consists of a series of convolutional and pooling layers. The convolutional layer has a brick-like structure on an 1D array of qubits, i.e. an m -qubit ($m < n$) trainable

convolutional ansatz is applied on all combinations of neighboring m qubits, mimicking the action of a kernel in CNN. It is a common practice to consider a two-qubit convolutional ansatz. All ansatzes within the same convolutional layer have same trainable parameters. In the pooling layer, a subset of qubits are measured and conditioned on each measurement outcome, a parametrized rotation is applied on each qubit of the complementary subset. Similar to convolutional layers the trainable parameters are shared within a pooling layer. All the measured qubits are then traced out, thus reducing the effective dimension of the system. In real quantum devices this is same as ignoring the qubits in the subsequent stages and considering only the dynamics on the remaining qubits. The sequence of convolutional and pooling layers is then repeated until a small fraction of qubits are left. In analogy to the fully-connected part of a CNN, one can then apply a general PQC on the remaining qubits at the end and finally measure them to obtain the prediction. The network is then trained using a suitable loss function and optimization algorithm. Compared to a generic deep quantum neural network, in this case due to progressive qubits reduction, the structure has a shallower depth of $\mathcal{O}(\log n)$, which is preferable for implementation in NISQ devices. This model of QCNN is also known to be free of barren plateau [41].

B. Label symmetry

For a classical or quantum dataset, a *label symmetry* exists if the class labels of all the data-points remain unchanged under a set of operations on those data-points. For example, in MNIST dataset the labels of digits $\{0, 1, 8\}$ are unchanged when the images are reflected, or the labels of a quantum state being entangled or non-entangled are invariant under local unitary transformations on a qubit. Let us consider a binary-classification task for a dataset $\mathcal{X}$ with data-points x_i and corresponding class labels $y_i \in \mathcal{Y}$ where $\mathcal{Y} = \{0, 1\}$. The function that maps the data-points to their labels is $f : \mathcal{X} \rightarrow \mathcal{Y}$. The task of QNN is to realize a quantum circuit f_θ that closely approximates f . A set of operations $\mathcal{G}$ form a label symmetry group of $\mathcal{X}$ if for all group elements $g \in \mathcal{G}$, the assigned labels remain unchanged, i.e.,

$$f(g(x_i)) = y_i = f(x_i), \quad \forall x_i \in \mathcal{X}, \forall g \in \mathcal{G}. \quad (1)$$

Note that the data points themselves may not be invariant under the group action, i.e., $g(x_i) \neq x_i$ in general. Thus, the parametric function f_θ used to classify $\mathcal{X}$ will be consistent with $\mathcal{G}$ if it predicts same label for inputs related by the symmetry operations. The QNNs respecting this constraint are called Equivariant QNNs (EQNNs). It was proved in [38, 53] that a PQC $\mathcal{U}_\theta$ and a measurement M will form an EQNN if and only if

$$[R(g), \mathcal{U}_\theta] = 0 \quad \text{and} \quad [R(g), M] = 0. \quad (2)$$

Thus if we define the *commutator space* of $R(g)$ to be the space of all operators that commute with $R(g)$, then $\mathcal{U}_\theta$ and M must belong to that commutator space. In the context




FIG. 1: In the center, a 4×4 image with pixel values $\{a, b, \dots, p\}$ indicated in the top left corner of each pixel and corresponding pixel positions in the bottom right corner. On the right, the image after vertical reflection. On the left, the image after rotation by angle $\pi/2$.

of EQCNN, $\mathcal{U}_\theta$ is composed of a series of convolutional and pooling layers. Thus, equivariance of $\mathcal{U}_\theta$ can be obtained by making sure that each convolution and pooling layer is equivariant with respect to $R(g)$.

C. EQCNN with permutation symmetry

The elements of a symmetry group can act either as *inner symmetry* or *outer symmetry* [38]. For the former, $R(g)$ can be decomposed as a tensor product of locally acting unitary operators. A well known example is the local unitary group $SU(2)$ acting on a system of n qubits, and its representation $R(g) = g^{\otimes n}$ which acts locally on each qubit, where $g \in SU(2)$. Another commonly occurring instance of inner symmetry is when classical images are encoded as quantum states using amplitude embedding, and the representation $R(g)$ of reflection or rotation group symmetry is a tensor product of Pauli- X matrices and qubit identity operators acting locally on different qubits. The architecture and performance of equivariant quantum neural networks for inner symmetries have been investigated in a number of works [37, 38, 44–46].

On the other hand, $R(g)$ for outer symmetry groups cannot be decomposed into locally acting components, an example being the Symmetric group S_n which manifests itself as a symmetry with respect to permuting the qubits. The physical operation of permuting two qubits is expressed mathematically using SWAP gate. One can show that, for every element g of S_n , R_g can be written as a product of SWAP gates, as we will exemplify in the next section. When a system of n qubits has full permutation symmetry, each qubit can be swapped with any other qubit. The fulfillment of equivariance condition Eq. (2) then requires that $\mathcal{U}_\theta$ must remain unchanged when the qubits are permuted in an arbitrary way. When using layers of single-qubit gates and two-qubit entangling gates to construct $\mathcal{U}_\theta$, permutation equivariance can be achieved if one uses single qubit gates with same parameters on all qubits, and entangling gates to connect all possible $\binom{n}{2}$ combinations of two qubits. In reference Ref. [42], the authors show that this architecture of S_n -equivariant QNN is devoid of barren plateau, and benchmark its performance for classification of

connected and non-connected graph states.

This structure of S_n -equivariant QNN is translationally symmetric on an array of qubits, thus can be used as the first convolutional layer of a permutation-equivariant QCNN. However, it is not straightforward to construct pooling layers in this case, as choosing a particular set of qubits to trace out inevitably breaks the full permutation symmetry. Additionally, the reduced representation on the remaining qubits after the first pooling layer is not a faithful representation. It was pointed out in Ref. [38] that one possible technique is to randomly trace out all combinations of $\binom{n}{n/2}$ qubits with equal probability to retain full permutation symmetry. In this work, we have used this strategy to construct equivariant pooling layers. Moreover, we show how to construct equivariant measurements for a probabilistic ansatz.

There exist other practically useful scenarios when the relevant labels of quantum states remain unchanged under the action of a subgroup of S_n . For example, when each pixel value of a classical image is encoded using a qubit (angle embedding), the operation of reflecting or rotating the image corresponds to permuting the qubits in a certain way, eliminating full freedom. In a way, this is an inner symmetry for which the group representation acts locally on a set of qubits which are allowed to permute within that set. In such cases, equivariant $\mathcal{U}_\theta$ can be achieved using less than $\binom{n}{2}$ two-qubit gates and a lower-depth circuit, as we will describe in the following subsections.

1. Reflection-equivariant QCNN

Let us consider classification of images whose labels remain invariant under a reflection about the vertical axis. There are two group elements here— the identity operation e which keeps the image unchanged, and the reflection operation r . They form the abstract group $\mathbb{Z}_2 \equiv \{e, r\}$. For angle embedding of classical images into quantum states, the symmetry representation of reflection group acts non-locally on certain groups of qubits through SWAP gates. For example, for a 4×4 image as in Fig. 1, the representation of the reflection group is,




FIG. 2: (a) The structure of reflection-equivariant QCNN for 16 qubits. The cyan and the orange boxes represent respectively the parametrized convolutional and the pooling ansatz. For the control qubits in pooling layer, we use half-filled circles to indicate that the pooling ansatz depends on both the control states $\{|0\rangle, |1\rangle\}$. (b) Reflection-equivariant QCNN using only nearest-neighbour connections in a 16-qubit quantum register with square-lattice architecture.

$$R^{\text{ref}} = \left\{ \mathbb{I}^{\otimes 16}, \text{SWAP}_{1,4} \text{SWAP}_{2,3} \text{SWAP}_{5,8} \text{SWAP}_{6,7} \text{SWAP}_{9,12} \text{SWAP}_{10,11} \text{SWAP}_{13,16} \text{SWAP}_{14,15} \right\} \quad (3)$$

where $\text{SWAP}_{i,j}$ denotes swap operation between pixel values at positions i and j of the image before reflection. Here, qubit pairs $\{\{1, 4\}, \{2, 3\} \dots\}$ constitute subsystems on each of which a SWAP gate acts locally.

Now, for an arbitrary order of embedding the pixel values in 1D array of qubits, one can check that applying convolutional ansatz in the typical brick-like structure is not equivariant with respect to the representation in Eq. (3). We demonstrate that by adhering to a particular pixel-to-qubit embed-

ding order, one can build equivariant convolutional ansatz for all layers by having the brick-like arrangement of two-qubit ansatz. By pixel-to-qubit embedding order, we imply a particular one-to-one correspondence between pixel values $\{a, b, c, \dots\}$ and the qubits indices $\{1, 2, 3, \dots\}$ which encodes them. This pixel-to-qubit embedding order and the resulting EQCNN have been presented in Fig. 2 for a 4×4 image. Here the qubit indices increase from 1 to 16 when moving from top to bottom of the 1D array. Due to this altered embedding order, the group representation changes to,

$$R^{\text{ref}} = \left\{ \mathbb{I}^{\otimes 16}, \text{SWAP}_{1,16} \text{SWAP}_{2,15} \text{SWAP}_{3,14} \text{SWAP}_{4,13} \text{SWAP}_{5,12} \text{SWAP}_{6,11} \text{SWAP}_{7,10} \text{SWAP}_{8,9} \right\}. \quad (4)$$

In the pooling layer, we trace out half of the qubits in a way so that the action of the reflection group has a faithful representation on the remaining qubits, which is the following after the first pooling layer,

$$R^{\text{ref}} = \left\{ \mathbb{I}^{\otimes 8}, \text{SWAP}_{5,12} \text{SWAP}_{6,11} \text{SWAP}_{7,10} \text{SWAP}_{8,9} \right\}. \quad (5)$$

Similarly, after the second pooling layer, the reduced representation is

$$R^{\text{ref}} = \left\{ \mathbb{I}^{\otimes 4}, \text{SWAP}_{6,11} \text{SWAP}_{7,10} \right\}, \quad (6)$$

and that after the third pooling layer is

$$R^{\text{ref}} = \left\{ \mathbb{I}^{\otimes 2}, \text{SWAP}_{7,10} \right\}. \quad (7)$$

The measurement is performed on the remaining two qubits $\{7, 10\}$.

This structure of EQCNN can be conveniently embedded on a quantum register with 2D square lattice architecture. We demonstrate this in Fig. 2(b). An advantage of 2D is that all the pooling ansatz are applied on nearest-neighbour qubits. However, to retain full translation symmetry of the convolutional layer, one needs to apply the convolutional ansatz on two qubits $\{a, d\}$ which are not nearest neighbours in this case. Square lattice architectures have been realized in NISQ devices such as Rigetti quantum hardware ANKAA-2 and Pasqal neutral atom quantum processing units.




FIG. 3: EQCNN for reflection and $\pi/2$ -rotation symmetry for 16 qubits. For non-nearest-neighbour applications of two-qubit convolutional ansatz in the first and third layer, we use same integers to clearly indicate qubits on which they apply within a sub-layer.

2. Reflection and $\frac{\pi}{2}$ -rotation equivariant QCNN

Next we consider images having label symmetry with respect to both reflection and a rotation by angle $\pi/2$. We

present the corresponding EQCNN in Fig. 3. There are two symmetry groups at play– the reflection symmetry group $\{e, r\}$, and the rotational symmetry group $\{e, r'\}$ where r' corresponds to $\pi/2$ -rotation. For the natural order of pixel-to-qubit embedding, the group representation for the latter is

$$R^{\text{rot}} = \left\{ \mathbb{I}^{\otimes 16}, \text{SWAP}_{1,16}\text{SWAP}_{2,15}\text{SWAP}_{3,14}\text{SWAP}_{4,13}\text{SWAP}_{5,12}\text{SWAP}_{6,11}\text{SWAP}_{7,10}\text{SWAP}_{8,9} \right\}, \quad (8)$$

as can be understood from Fig. 1. Note that Eq. (8) is the same as the representation in Eq. (4). Hence in this case, rotational equivariance is obtained from the brick-like structure of convolutional layer for the natural order of pixel-to-qubit embedding. However, to ensure reflection equivariance for this embedding, an additional layer of convolutional ansatz needs to be applied but in a different fashion. In detail, the two-qubit

convolutional ansatz is applied on every qubit and its 7th nearest neighbour when counting the qubits from top to bottom in Fig. 3. The reasoning behind this construction can be understood by observing how the nearest neighbour qubits are rearranged under reflection. The construction of first pooling layer ensures that the representations of both the groups on the remaining qubits are faithful. These reduced representations are

$$R^{\text{ref}} = \left\{ \mathbb{I}^{\otimes 8}, \text{SWAP}_{5,8}\text{SWAP}_{6,7}\text{SWAP}_{9,12}\text{SWAP}_{10,11} \right\}, \quad (9)$$

$$R^{\text{rot}} = \left\{ \mathbb{I}^{\otimes 8}, \text{SWAP}_{5,12}\text{SWAP}_{6,11}\text{SWAP}_{7,10}\text{SWAP}_{8,9} \right\}. \quad (10)$$

In the second convolutional layer, the brick-like structure alone is sufficient to ensure equivariance with respect to both the symmetries. The reduced representations after the second pooling layer are,

$$R^{\text{ref}} = \left\{ \mathbb{I}^{\otimes 4}, \text{SWAP}_{5,8}\text{SWAP}_{9,12} \right\}, \quad (11)$$

$$R^{\text{rot}} = \left\{ \mathbb{I}^{\otimes 4}, \text{SWAP}_{5,12}\text{SWAP}_{8,9} \right\}. \quad (12)$$

The third convolutional layer is constructed in the same way as the second convolutional layer. In the third pooling layer, it is impossible to trace out any two qubits, maintaining equivariance with respect to both the symmetry groups. Hence, we use a pooling layer for which the rotational equivariance is broken but the reflection equivariance is protected. Lastly, the remaining two qubits $\{5, 8\}$ are measured.

Note that, one can obtain the EQCNN with respect to $\pi/2$ -




FIG. 4: (a) A schematic diagram showing all possible QCNNs and the qubit indices on which they act for an S_4 -equivariant QCNN. Each arrow indicates a possible pooling operation. The number of different pooling operations in a layer is shown in the vertical column on the left. We highlight two examples of these QCNNs with red and blue arrows. (b) A circuit representation of the red and blue QCNNs and corresponding measurements. C_1 and C_2 are respectively the first and second convolutional layer.

rotation symmetry alone by removing the additional convolutional layer in the first layer, and changing the last pooling ansatz in a way to protect rotational equivariance. As a result, the final measurement will be on a different set of qubits.

3. S_n -equivariant QCNN

When a system has full permutation symmetry, we use the S_n -equivariant QNN with all-to-all two-qubit gate connections as the first convolutional layer. For the pooling operation, let us consider a scenario in which no explicit pooling ansatz is applied, but half of the qubits are traced out successively at each pooling layer. In the first pooling layer, we randomly trace out any $n/2$ qubits. This can be done in $\binom{n}{n/2}$ possible ways, each occurring with equal probability to retain the permutation equivariance. The second convolutional layer has the same structure as the first convolutional layer, however it is applied on the remaining $n/2$ qubits. In the second pooling layer, we randomly trace out $n/4$ qubits from $n/2$ qubits in $\binom{n/2}{n/4}$ possible ways, and apply the third convolutional layer on the remaining $n/4$ qubits. We keep adding convolutional and pooling layers until only one qubit is left. Another way to envision this is to think that there are in total

$$P = \binom{n}{n/2} \times \binom{n/2}{n/4} \times \binom{n/4}{n/8} \times \dots \times \binom{2}{1} \quad (13)$$

number of possible QCNNs, which are applied on the input state with equal probability. For each of these QCNNs, the last remaining qubit is measured to obtain the loss function. We schematically depict the idea discussed above in Fig. 4 for a simple case of $n = 4$.

Now we build a quantum circuit that can exactly realize the probabilistic nature of the EQCNN described above. We show this in Fig. 5. Additional to the target register encoding the input state ρ , there is an auxiliary quantum register A_1 with $\lceil \log P \rceil = 4$ qubits, which is initialized in an equal superposition $|\Psi_{A_1}\rangle$ of the first $P = 12$ orthogonal basis states. We

apply the first convolutional layer C_1 on ρ . Since no trainable pooling ansatz is present, tracing out $n/2$ randomly chosen qubits in the first pooling layer is equivalent to applying the second convolutional layer C_2 on all possible $\binom{4}{2} = 6$ pairs of qubits. To ensure that they are applied at random with equal probability, we apply them controlled on the states of the first three qubits in A_1 . In the next and final pooling layer, randomly tracing out any one qubit from a qubit pair is equivalent to randomly choosing a qubit and measuring it. For this, we add another single-qubit auxiliary register A_2 initialized in state $|0\rangle$. Now for each qubit pair in the target register, controlled on the last qubit of A_1 being in state $|0\rangle$ or $|1\rangle$, we swap either of the qubits with A_2 . Finally, we measure Pauli-Z operator on A_2 to obtain the loss function. A detailed mathematical description of how this circuit works is discussed in the Appendix.

A drawback of the above circuit is that with increasing n , P increases very fast and soon becomes exponential in n , resulting in an exponentially deep circuit. This can be avoided if instead of tracing out half of the qubits from previous layer, one chooses to trace out a constant number m of qubits at each pooling layer, and repeats this for $k < \frac{n}{m}$ times such that P is polynomial in n . At the end of k^{th} layer, measurements are performed on all the remaining qubits. This strategy remains valid since tracing out exactly half of the qubits is a convention rather than a strict requirement for QCNN, although the logarithmic depth may not hold depending on the relation between m and n . We leave it for future research works to investigate the performance of this polynomially deep EQCNN.

As mentioned in ref. [38], the technique of probabilistically applying different unitary ansatz on the input states can also be used to design a dropout mechanism in a general quantum neural network. A few works exist in literature that proposes quantum dropout by randomly dropping single and two-qubit gates [54, 55] as well as the qubits themselves [10], whereas our strategy matches well with that discussed in ref. [56]. Particularly, ours is a special case in which half of the qubits from the previous layer are randomly chosen and dropped along




FIG. 5: The circuit for realizing EQCNN for symmetric group S_4 on the input state ρ . In the controlled gates, the circles with no fillings and black fillings represent states $|0\rangle$ and $|1\rangle$ respectively.




FIG. 6: Samples of Fashion MNIST images corresponding to class 0 and class 8 sized down to 4×4 pixels.

with all gates acting on them on subsequent layers. However, it also implies dropping half of the trainable parameters at each layer, which greatly reduces the expressibility and can result in poor trainability of the network. Instead, one may choose to drop a small constant number of m qubits at each successive layer, as discussed also in the previous paragraph.

III. RESULTS

We compare the classification accuracy of equivariant and non-equivariant QCNs for the symmetry groups discussed in the last section. We build and train all the networks using PennyLane quantum simulator [57].

For reflection symmetry alone, we use class 0 (tshirt) and class 8 (handbag) of Fashion MNIST dataset, whereas for images symmetric with respect to reflection and $\pi/2$ -rotation, we use class 0 and 1 of MNIST dataset. The original images are 28×28 pixels, we downsize them to 4×4 pixels to embed the pixel values into 16 qubits using R_Y rotation. Some samples of the downsized images for Fashion MNIST dataset are presented in Fig. 6. As the convolutional and pooling ansatz, we use the circuits in Fig. 7(a) and 7(b) respectively, where $\text{Rot}(\theta_1, \theta_2, \theta_3) = R_Z(\theta_3)R_Y(\theta_2)R_Z(\theta_1)$. One can choose any other set of single-qubit and two-qubit gates without loss of generality. For building EQCNs we follow the

embeddings and architectures in Fig. 2 and Fig. 3, while for non-equivariant QCNs we use standard pixel-to-qubit embedding order and the architectures in Ref. [51]. We train the models using Nesterov moment optimizer with learning rate 0.005. As loss function, we use binary cross-entropy. We obtain the average test-set classification accuracy over 10 identical runs, each with random initial parameters and a batch size of 32. The behaviour of mean test-set accuracy with increasing number of training iterations is presented in Fig. 8. For fashion MNIST, EQCNN has a higher classification accuracy than non-equivariant QCNN. For MNIST images, considering the $\pi/2$ -rotational symmetry alone results in a higher accuracy than the non-equivariant QCNN. However, taking into account both reflection and $\pi/2$ -rotational symmetry reduces the accuracy significantly. This can be due to the highly constrained structure of EQCNN under both the symmetries, for which the expressibility degrades significantly.

We apply the S_n -equivariant QCNN for classification of connected and non-connected graphs with 4 vertices. The graphs are randomly generated using Erdős-Rényi model with an edge probability of 45%. The quantum graph states are generated from the graphs using the standard method from [58, 59], and are embedded into a four-qubit quantum register using amplitude embedding. The property of being a connected or non-connected graph remains unchanged under arbitrary permutations of the vertices, in this case the qubits. From all possible graph states with 4 vertices, we choose a fraction for training and the other part for testing. We use two different instances of this partitioning- for *case 1* we have 45 data-points for training and 18 for testing, while for *case 2* we have 52 and 11 data-points respectively for training and testing. In both cases, the training data is well balanced between the two classes, while the test data has a large fraction of connected graphs. We augment the training and test set with multiple copies of the distinct data-points. As the convolutional ansatz, we use the S_n -equivariant ansatz shown in Fig. 7(c) in the first layer, and its two-qubit analog in the second layer. The choice of the gates is again arbitrary. In addition to EQCNN, we also test the performance of S_n -




FIG. 7: (a) The convolutional ansatz and (b) the pooling ansatz used to construct EQCNN and non-equivariant QCNN for reflection and $\pi/2$ -rotation symmetry. (c) The convolutional ansatz used in S_n -equivariant QCNN.




FIG. 8: Comparison of EQCNN and non-equivariant QCNN for binary classification of (a) classes 0 and 8 of Fashion MNIST dataset and (b) classes 0 and 1 of MNIST dataset. All quantities are averaged over 10 training iterations with random initial parameters.

equivariant QNN without pooling layer, for which we replace the second convolutional layer by another layer of the four-qubit ansatz in Fig. 7(c). In this case we measure Pauli-Z operator on all the qubits to obtain the loss function. Note that the number of trainable parameters is same for the EQCNN and EQNN. For non-equivariant QCNN, we use the architecture of QCNN in Ref. [51] with Fig. 7(a) and 7(b) as convolutional and pooling ansatz. All models are trained using Adam optimizer with a learning rate 0.01, and we use mean-squared error as loss function. We obtain the mean training loss, training accuracy and test accuracy from 10 randomly initialized training schemes with a batch size 10. The results are presented in Fig. 9.

On average, we find that the S_n -EQNN has a good generalization behaviour, and the corresponding test accuracy is better than the other two models for case 1 when the train-test ratio is 5 : 3 approximately. The performance of S_n -EQNN does not differ much for case 1 and case 2. On the other hand, the S_n -EQCNN overfits the training data in case 1. However, when the train-test ratio is approximately 5 : 1 in case 2 and there are more data-points to learn, the performance of the EQCNN improves and it shows higher average test accuracy

than S_n -equivariant QNN. In this case, out of the 10 instances of initial parameters, there are more cases for which the maximum test accuracy 1 is obtained for EQCNN compared to EQNN. In Fig. 10 we show the training and test accuracy for one such instance. The non-equivariant QCNN has a much lower test accuracy than both the equivariant architectures in both cases.

IV. DISCUSSIONS AND CONCLUSION

In this work, we proposed a pixel-to-qubit embedding order that facilitates applying the quantum convolutional ansatz in an equivariant manner for angle-embedded classical images with reflection and rotational symmetry of image labels. The resulting EQCNNs show better classification accuracy than non-equivariant QCNNs. Though angle embedding requires as many qubits as the number of pixels compared to the logarithmically less number of qubits in amplitude embedding, the latter needs an exponentially deep circuit to encode an arbitrary quantum superposition of n qubits, thus subsuming the advantage due to logarithmically shallow circuit of QCNN.




FIG. 9: Comparison of training set loss and accuracy, and test set classification accuracy of EQCNN, EQNN, and non-equivariant QCNN for classification of connected and non-connected graphs with 4 vertices. All quantities are averaged over 10 training iterations with random initial parameters. (a) The training set and test set has respectively 45 and 18 distinct data-points. (b) The training set and test set has respectively 52 and 11 distinct data-points.




FIG. 10: An instance of initial parameters for which S_n -EQCNN performs better than S_n -EQNN when the train-test split has 5 : 1 ratio.

Compared to that, angle embedding requires a circuit of constant depth, since each qubit can be manipulated individually in parallel. With this advantage, and the fast growing dimen-

sion of NISQ devices, angle embedding remains a more potent choice for quantum embedding of classical data. Moreover, one can also pre-process a large classical image to reduce its dimension using principal component analysis (PCA) or a classical autoencoder, before angle-embedding it into a smaller number of available qubits. In this case however, the pre-processing must preserve the label symmetry, in order to use an EQCNN for classification. It is to be noted that, though our embedding trick works for small subgroups of S_n , it will become increasingly difficult to design a translationally symmetric equivariant convolutional ansatz for bigger permutation groups, e.g. $p4m$ symmetry group, which again arises frequently in highly symmetric classical images. In such cases, one may choose to take into account fewer symmetries. This is supported by our numerical data showing a better performance compared to equivariant circuit which takes into account all existing symmetries. One can also use a more flexible ansatz that weakly breaks the symmetries [45].

Our construction of S_n -equivariant QCNN proposed also a probabilistic picture of QNN, which has not been explored in detail in literature yet. We use graph states as dataset to analyze the performance of the S_n -equivariant EQCNN, which shows a better performance than non-equivariant QCNN.

When sufficient training data are used, this EQCNN performs better than S_n -equivariant QNN on average. This probabilistic application of S_n -equivariant QCNN can be viewed as a special type of dropout mechanism in QNN [10, 54–56], in analogy to dropout in CNN to prevent overfitting [60]. Particularly in our case, the important difference with the conventional dropout is that, here no trainable parameters are dropped, rather some of the qubits on which the corresponding gates apply. One can also apply with equal probability the equivariant ansatz for every symmetry group when there exist more than one symmetry group, as also discussed in the last paragraph.

It should be mentioned that there exist alternative approaches for designing a quantum convolutional neural network in quantum systems [13, 39, 61]. Particularly in references [39, 61], the authors conceptualize the qubit permutation symmetry as the quantum analog of translational symmetry in classical images. They propose as ‘convolutional quantum ansatz’ a permutation-equivariant ansatz which is in turn constructed by employing global $SU(d)$ -equivariant Hamiltonians, since by Schur’s lemma the representation theory of S_n can be described using the representation theory of $SU(d)$, and vice versa. As a result, their architecture can be used to find the ground states of many-body Hamiltonians that exhibit $SU(d)$ symmetry. This is in contrast to our work which proposes a permutation-equivariant QCNN that manifests the translation symmetry in a different way [11].

Data classification is an ubiquitous task in a plethora of everyday applications. Classical state-of-the-art classifiers are supremely successful, yet the potency of quantum classifiers are worth the investigation. The quantum convolutional neural networks are proved to be efficient classifiers. Thus, our construction of equivariant QCNNs for spatial symmetry of images is a significant addition to the growing literature. The S_n -equivariant QCNN proposed in this work is extremely relevant due to the frequent presence of permutation symmetry in nature. We also believe that the stochastic nature of the variational quantum ansatz is an interesting direction to be pursued for future research.

ACKNOWLEDGMENTS

This work was supported by the European Commission’s Horizon Europe Framework Programme under the Research

and Innovation Action GA n. 101070546–MUQUABIS, by the European Union’s Horizon 2020 research and innovation programme under FET-OPEN GA n. 828946–PATHOS, by the European Defence Agency under the project Q-LAMPS Contract No B PRJ- RT-989, and by the MUR Progetti di Ricerca di Rilevante Interesse Nazionale (PRIN) Bando 2022 - project n. 20227HSE83 – ThAI-MIA funded by the European Union - Next Generation EU.

Appendix A: S_n -equivariant channel and measurements

Denoting the input state as ρ , the state after the first convolutional layer is

$$\rho_1 = C_1 \rho C_1^\dagger, \quad (\text{A1})$$

where C_1 is the unitary corresponding to the first convolutional layer.

The initial state in the auxiliary register A_1 is

$$\begin{aligned} |\Psi_{A_1}\rangle = & \frac{1}{\sqrt{12}}(|0000\rangle + |0001\rangle + |0010\rangle + |0011\rangle + |0100\rangle \\ & + |0101\rangle + |0110\rangle + |0111\rangle + |1000\rangle + |1001\rangle \\ & + |1010\rangle + |1011\rangle). \end{aligned} \quad (\text{A2})$$

After the controlled application of second convolutional layer on different pairs of qubits, the input state becomes,

$$\begin{aligned} \rho_2 = & \frac{1}{6}(C_2^{12} \rho_1 C_2^{12\dagger} + C_2^{13} \rho_1 C_2^{13\dagger} + C_2^{14} \rho_1 C_2^{14\dagger} \\ & + C_2^{23} \rho_1 C_2^{23\dagger} + C_2^{24} \rho_1 C_2^{24\dagger} + C_2^{34} \rho_1 C_2^{34\dagger}), \end{aligned} \quad (\text{A3})$$

where C_2^{ij} is the second convolutional layer acting on the qubit-pair (i, j) .

At this stage, the auxiliary register A_2 is in the state $|0\rangle$, and we denote the corresponding qubit with index 0. After application of controlled swap gates, the state of the joint system consisting of our input state register and auxiliary register A_2 is,

$$\begin{aligned} \rho'_3 = & \frac{1}{12}(S_{01}(C_2^{12} \rho_1 C_2^{12\dagger} \otimes |0\rangle\langle 0|)S_{01}^\dagger + S_{02}(C_2^{12} \rho_1 C_2^{12\dagger} \otimes |0\rangle\langle 0|)S_{02}^\dagger + S_{01}(C_2^{13} \rho_1 C_2^{13\dagger} \otimes |0\rangle\langle 0|)S_{01}^\dagger \\ & + S_{03}(C_2^{13} \rho_1 C_2^{13\dagger} \otimes |0\rangle\langle 0|)S_{03}^\dagger + S_{01}(C_2^{14} \rho_1 C_2^{14\dagger} \otimes |0\rangle\langle 0|)S_{01}^\dagger + S_{04}(C_2^{14} \rho_1 C_2^{14\dagger} \otimes |0\rangle\langle 0|)S_{04}^\dagger \\ & + S_{02}(C_2^{23} \rho_1 C_2^{23\dagger} \otimes |0\rangle\langle 0|)S_{02}^\dagger + S_{03}(C_2^{23} \rho_1 C_2^{23\dagger} \otimes |0\rangle\langle 0|)S_{03}^\dagger + S_{02}(C_2^{24} \rho_1 C_2^{24\dagger} \otimes |0\rangle\langle 0|)S_{02}^\dagger \\ & + S_{04}(C_2^{24} \rho_1 C_2^{24\dagger} \otimes |0\rangle\langle 0|)S_{04}^\dagger + S_{03}(C_2^{34} \rho_1 C_2^{34\dagger} \otimes |0\rangle\langle 0|)S_{03}^\dagger + S_{04}(C_2^{34} \rho_1 C_2^{34\dagger} \otimes |0\rangle\langle 0|)S_{04}^\dagger). \end{aligned}$$

The expectation value of the measurement on A_2 is,

$$m = \text{Tr}(\rho'_3 M'_0), \quad (\text{A4})$$

where $M'_0 = \sigma_0^\dagger \otimes \mathbb{I}_1 \otimes \mathbb{I}_2 \otimes \mathbb{I}_3 \otimes \mathbb{I}_4$.

Now, using the cyclic property of trace, we can write the

following for an arbitrary term in Eq. (A4),

$$\begin{aligned}
 m_i^{ij} &= \text{Tr} \left((S_{0i} (C_2^{ij} \rho_1 C_2^{ij\dagger} \otimes |0\rangle\langle 0|) S_{0i}^\dagger) M'_0 \right) \\
 &= \text{Tr} \left((S_{0i}^\dagger M'_0 S_{0i}) (C_2^{ij} \rho_1 C_2^{ij\dagger} \otimes |0\rangle\langle 0|) \right) \\
 &= \text{Tr} (M'_i (C_2^{ij} \rho_1 C_2^{ij\dagger} \otimes |0\rangle\langle 0|)) \\
 &= \text{Tr} (M_i C_2^{ij} \rho_1 C_2^{ij\dagger}) \cdot \text{Tr} (\mathbb{I}_0 |0\rangle\langle 0|) \\
 &= \text{Tr} (M_i C_2^{ij} \rho_1 C_2^{ij\dagger}),
 \end{aligned}$$

where M_i ($i \neq 0$) denotes a Pauli-Z measurement on the i^{th} qubit. In this way, we design a measurement outcome which is the weighted sum of measurement outcomes from all equally probable QCNNs.

-
- [1] M. Schuld, I. Sinayskiy, and F. Petruccione, An introduction to quantum machine learning, *Contemporary Physics* **56**, 172 (2015), <https://doi.org/10.1080/00107514.2014.964942>.
 - [2] J. Biamonte, P. Wittek, N. Pancotti, P. Rebentrost, N. Wiebe, and S. Lloyd, Quantum machine learning, *Nature* **549**, 195 (2017).
 - [3] M. Benedetti, E. Lloyd, S. Sack, and M. Fiorentini, Parametrized quantum circuits as machine learning models, *Quantum Science and Technology* **4**, 043001 (2019).
 - [4] A. Perdomo-Ortiz, M. Benedetti, J. Realpe-Gómez, and R. Biswas, Opportunities and challenges for quantum-assisted machine learning in near-term quantum computers, *Quantum Science and Technology* **3**, 030502 (2018).
 - [5] M. Schuld and N. Killoran, Is quantum advantage the right goal for quantum machine learning?, *PRX Quantum* **3**, 030101 (2022).
 - [6] A. Peruzzo, J. McClean, P. Shadbolt, M.-H. Yung, X.-Q. Zhou, P. J. Love, A. Aspuru-Guzik, and J. L. O'Brien, A variational eigenvalue solver on a photonic quantum processor, *Nature Communications* **5**, 4213 (2014).
 - [7] J. R. McClean, J. Romero, R. Babbush, and A. Aspuru-Guzik, The theory of variational hybrid quantum-classical algorithms, *New Journal of Physics* **18**, 023023 (2016).
 - [8] J. Romero, J. P. Olson, and A. Aspuru-Guzik, Quantum autoencoders for efficient compression of quantum data, *Quantum Science and Technology* **2**, 045001 (2017).
 - [9] K. Mitarai, M. Negoro, M. Kitagawa, and K. Fujii, Quantum circuit learning, *Phys. Rev. A* **98**, 032309 (2018).
 - [10] M. Schuld, A. Bocharov, K. M. Svore, and N. Wiebe, Circuit-centric quantum classifiers, *Phys. Rev. A* **101**, 032308 (2020).
 - [11] I. Cong, S. Choi, and M. D. Lukin, Quantum convolutional neural networks, *Nature Physics* **15**, 1273 (2019).
 - [12] Y. Li, R.-G. Zhou, R. Xu, J. Luo, and W. Hu, A quantum deep convolutional neural network for image recognition, *Quantum Science and Technology* **5**, 044003 (2020).
 - [13] M. Henderson, S. Shakyia, S. Pradhan, and T. Cook, Quantum convolutional neural networks: powering image recognition with quantum circuits, *Quantum Machine Intelligence* **2**, 2 (2020).
 - [14] P. Braccia, F. Caruso, and L. Banchi, How to enhance quantum generative adversarial learning of noisy information, *New Journal of Physics* **23**, 053024 (2021).
 - [15] P. Braccia, L. Banchi, and F. Caruso, Quantum noise sensing by generating fake noise, *Phys. Rev. Appl.* **17**, 024002 (2022).
 - [16] M. S. Rudolph, N. B. Toussaint, A. Katabarwa, S. Johri, B. Peropadre, and A. Perdomo-Ortiz, Generation of high-resolution handwritten digits with an ion-trap quantum computer, *Phys. Rev. X* **12**, 031010 (2022).
 - [17] A. O. Boyle and R. Nikandish, A hybrid quantum-classical generative adversarial network for near-term quantum processors (2023), [arXiv:2307.03269 \[quant-ph\]](https://arxiv.org/abs/2307.03269).
 - [18] S. L. Tsang, M. T. West, S. M. Erfani, and M. Usman, Hybrid quantum-classical generative adversarial network for high-resolution image generation, *IEEE Transactions on Quantum Engineering* **4**, 1 (2023).
 - [19] N.-R. Zhou, T.-F. Zhang, X.-W. Xie, and J.-Y. Wu, Hybrid quantum-classical generative adversarial networks for image generation via learning discrete distribution, *Signal Processing: Image Communication* **110**, 116891 (2023).
 - [20] M. Parigi, S. Martina, and F. Caruso, Quantum-noise-driven generative diffusion models (2023), [arXiv:2308.12013 \[quant-ph\]](https://arxiv.org/abs/2308.12013).
 - [21] B. Zhang, P. Xu, X. Chen, and Q. Zhuang, Generative quantum machine learning via denoising diffusion probabilistic models, *Phys. Rev. Lett.* **132**, 100602 (2024).
 - [22] A. Cacioppo, L. Colantonio, S. Bordoni, and S. Giagu, Quantum diffusion models (2023), [arXiv:2311.15444 \[quant-ph\]](https://arxiv.org/abs/2311.15444).
 - [23] P. Rebentrost, M. Mohseni, and S. Lloyd, Quantum support vector machine for big data classification, *Phys. Rev. Lett.* **113**, 130503 (2014).
 - [24] Y. Liu, S. Arunachalam, and K. Temme, A rigorous and robust quantum speed-up in supervised machine learning, *Nature Physics* **17**, 1013 (2021).
 - [25] A. Abbas, D. Sutter, C. Zoufal, A. Lucchi, A. Figalli, and S. Woerner, The power of quantum neural networks, *Nature Computational Science* **1**, 403 (2021).
 - [26] 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), <https://www.science.org/doi/pdf/10.1126/science.abn7293>.
 - [27] 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, *Nature Communications* **13**, 4919 (2022).
- [28] J. R. McClean, S. Boixo, V. N. Smelyanskiy, R. Babbush, and H. Neven, Barren plateaus in quantum neural network training landscapes, *Nature Communications* **9**, 4812 (2018).
- [29] Z. Holmes, K. Sharma, M. Cerezo, and P. J. Coles, Connecting ansatz expressibility to gradient magnitudes and barren plateaus, *PRX Quantum* **3**, 010313 (2022).
- [30] M. Cerezo, A. Sone, T. Volkoff, L. Cincio, and P. J. Coles, Cost function dependent barren plateaus in shallow parametrized quantum circuits, *Nature Communications* **12**, 1791 (2021).
- [31] S. Wang, E. Fontana, M. Cerezo, K. Sharma, A. Sone, L. Cincio, and P. J. Coles, Noise-induced barren plateaus in variational quantum algorithms, *Nature Communications* **12**, 6961 (2021).
- [32] S. Sim, P. D. Johnson, and A. Aspuru-Guzik, Expressibility and entangling capability of parameterized quantum circuits for hybrid quantum-classical algorithms, *Advanced Quantum Technologies* **2**, 1900070 (2019), <https://onlinelibrary.wiley.com/doi/pdf/10.1002/qute.201900070>.
- [33] M. M. Bronstein, J. Bruna, T. Cohen, and P. Veličković, Geometric deep learning: Grids, groups, graphs, geodesics, and gauges (2021), [arXiv:2104.13478 \[cs.LG\]](https://arxiv.org/abs/2104.13478).
- [34] 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).
- [35] P. Mernyei, K. Meichanetzidis, and İsmail İlkan Ceylan, Equivariant quantum graph circuits (2022), [arXiv:2112.05261 \[cs.LG\]](https://arxiv.org/abs/2112.05261).
- [36] A. Skolik, M. Cattelan, S. Yarkoni, T. Bäck, and V. Dunjko, Equivariant quantum circuits for learning on weighted graphs, *npj Quantum Information* **9**, 47 (2023).
- [37] 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).
- [38] Q. T. Nguyen, L. Schatzki, P. Braccia, M. Ragone, P. J. Coles, F. Sauvage, M. Larocca, and M. Cerezo, Theory for equivariant quantum neural networks (2022), [arXiv:2210.08566 \[quant-ph\]](https://arxiv.org/abs/2210.08566).
- [39] 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).
- [40] R. D. P. East, G. Alonso-Linaje, and C.-Y. Park, All you need is spin: $Su(2)$ equivariant variational quantum circuits based on spin networks (2023), [arXiv:2309.07250 \[quant-ph\]](https://arxiv.org/abs/2309.07250).
- [41] 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).
- [42] L. Schatzki, M. Larocca, Q. T. Nguyen, F. Sauvage, and M. Cerezo, Theoretical guarantees for permutation-equivariant quantum neural networks, *npj Quantum Information* **10**, 12 (2024).
- [43] M. T. West, J. Heredge, M. Sevier, and M. Usman, Provably trainable rotationally equivariant quantum machine learning (2023), [arXiv:2311.05873 \[quant-ph\]](https://arxiv.org/abs/2311.05873).
- [44] M. T. West, M. Sevier, and M. Usman, Reflection equivariant quantum neural networks for enhanced image classification, *Machine Learning: Science and Technology* **4**, 035027 (2023).
- [45] S. Y. Chang, M. Grossi, B. Le Saux, and S. Vallecorsa, Approximately equivariant quantum neural network for p4m group symmetries in images, in *2023 IEEE International Conference on Quantum Computing and Engineering (QCE)*, Vol. 01 (2023) pp. 229–235.
- [46] S. Das, S. Martina, and F. Caruso, The role of data embedding in equivariant quantum convolutional neural networks (2024), [arXiv:2312.13250 \[quant-ph\]](https://arxiv.org/abs/2312.13250).
- [47] J. Herrmann, S. M. Llima, A. Remm, P. Zapletal, N. A. McMahon, C. Scarato, F. Swiadek, C. K. Andersen, C. Hellings, S. Krinner, N. Lacroix, S. Lazar, M. Kerschbaum, D. C. Zanuz, G. J. Norris, M. J. Hartmann, A. Wallraff, and C. Eichler, Realizing quantum convolutional neural networks on a superconducting quantum processor to recognize quantum phases, *Nature Communications* **13**, 4144 (2022).
- [48] Y.-J. Liu, A. Smith, M. Knap, and F. Pollmann, Model-independent learning of quantum phases of matter with quantum convolutional neural networks, *Phys. Rev. Lett.* **130**, 220603 (2023).
- [49] C. Umeano, A. E. Paine, V. E. Elfving, and O. Kyriienko, What can we learn from quantum convolutional neural networks? (2023), [arXiv:2308.16664 \[quant-ph\]](https://arxiv.org/abs/2308.16664).
- [50] S. Monaco, O. Kiss, A. Mandarino, S. Vallecorsa, and M. Grossi, Quantum phase detection generalization from marginal quantum neural network models, *Phys. Rev. B* **107**, L081105 (2023).
- [51] T. Hur, L. Kim, and D. K. Park, Quantum convolutional neural network for classical data classification, *Quantum Machine Intelligence* **4**, 3 (2022).
- [52] G. Vidal, Class of quantum many-body states that can be efficiently simulated, *Phys. Rev. Lett.* **101**, 110501 (2008).
- [53] M. Ragone, P. Braccia, Q. T. Nguyen, L. Schatzki, P. J. Coles, F. Sauvage, M. Larocca, and M. Cerezo, Representation theory for geometric quantum machine learning (2023), [arXiv:2210.07980 \[quant-ph\]](https://arxiv.org/abs/2210.07980).
- [54] M. Kobayashi, K. Nakaji, and N. Yamamoto, Overfitting in quantum machine learning and entangling dropout, *Quantum Machine Intelligence* **4**, 30 (2022).
- [55] F. Scala, A. Ceschini, M. Panella, and D. Gerace, A general approach to dropout in quantum neural networks, *Advanced Quantum Technologies* **n/a**, 2300220, <https://onlinelibrary.wiley.com/doi/pdf/10.1002/qute.202300220>.
- [56] G. Verdon, J. Pye, and M. Broughton, A universal training algorithm for quantum deep learning (2018), [arXiv:1806.09729 \[quant-ph\]](https://arxiv.org/abs/1806.09729).
- [57] V. Bergholm, J. Izaac, M. Schuld, C. Gogolin, S. Ahmed, V. Ajith, M. S. Alam, G. Alonso-Linaje, B. AkashNarayanan, A. Asadi, J. M. Arrazola, U. Azad, S. Banning, C. Blank, T. R. Bromley, B. A. Cordier, J. Ceroni, A. Delgado, O. D. Matteo, A. Dusko, T. Garg, D. Guala, A. Hayes, R. Hill, A. Ijaz, T. Isaacsson, D. Ittah, S. Jahangiri, P. Jain, E. Jiang, A. Khandelwal, K. Kottmann, R. A. Lang, C. Lee, T. Loke, A. Lowe, K. McKiernan, J. J. Meyer, J. A. Montañez-Barrera, R. Moyard, Z. Niu, L. J. O’Riordan, S. Oud, A. Panigrahi, C.-Y. Park, D. Polatajko, N. Quesada, C. Roberts, N. Sá, I. Schoch, B. Shi, S. Shu, S. Sim, A. Singh, I. Strandberg, J. Soni, A. Száva, S. Thabet, R. A. Vargas-Hernández, T. Vincent, N. Vitucci, M. Weber, D. Wierichs, R. Wiersema, M. Willmann, V. Wong, S. Zhang, and N. Killoran, PennyLane: Automatic differentiation of hybrid quantum-classical computations (2022), [arXiv:1811.04968 \[quant-ph\]](https://arxiv.org/abs/1811.04968).
- [58] R. Raussendorf, D. E. Browne, and H. J. Briegel, Measurement-based quantum computation on cluster states, *Phys. Rev. A* **68**, 022312 (2003).
- [59] M. Hein, J. Eisert, and H. J. Briegel, Multiparty entanglement in graph states, *Phys. Rev. A* **69**, 062311 (2004).
- [60] N. Srivastava, G. Hinton, A. Krizhevsky, I. Sutskever, and R. Salakhutdinov, Dropout: A simple way to prevent neural networks from overfitting, *Journal of Machine Learning Research* **15**, 1929 (2014).
- [61] H. Zheng, Z. Li, J. Liu, S. Strelchuk, and R. Kondor, On

the super-exponential quantum speedup of equivariant quan-

tum machine learning algorithms with $\text{su}(d)$ symmetry (2022),
[arXiv:2207.07250 \[quant-ph\]](#).