






Splitting and parallelizing of quantum convolutional neural networks for learning translationally symmetric data

Koki Chinzei *, Quoc Hoan Tran , Kazunori Maruyama , Hiroataka Oshima , and Shintaro Sato
Quantum Laboratory, Fujitsu Research, Fujitsu Limited, 4-1-1 Kawasaki, Kanagawa 211-8588, Japan

 (Received 20 June 2023; accepted 23 January 2024; published 11 April 2024)

The quantum convolutional neural network (QCNN) is a promising quantum machine learning (QML) model that is expected to achieve quantum advantages in classically intractable problems. However, the QCNN requires a large number of measurements for data learning, limiting its practical applications in large-scale problems. To alleviate this requirement, we propose an architecture called split-parallelizing QCNN (sp-QCNN), which exploits the prior knowledge of quantum data to design an efficient model. This architecture draws inspiration from geometric quantum machine learning and targets translationally symmetric quantum data commonly encountered in physics and quantum computing science. By splitting the quantum circuit based on translational symmetry, the sp-QCNN can substantially parallelize the conventional QCNN without increasing the number of qubits and improve the measurement efficiency by an order of the number of qubits. To demonstrate its effectiveness, we apply the sp-QCNN to a quantum phase recognition task and show that it can achieve comparable classification accuracy to the conventional QCNN while considerably reducing the measurement resources required. Due to its high measurement efficiency, the sp-QCNN can mitigate statistical errors in estimating the gradient of the loss function, thereby accelerating the learning process. These results open up possibilities for incorporating the prior data knowledge into the efficient design of QML models, leading to practical quantum advantages.

DOI: [10.1103/PhysRevResearch.6.023042](https://doi.org/10.1103/PhysRevResearch.6.023042)

I. INTRODUCTION

Quantum computing is an innovative technology that is expected to solve classically intractable problems and open up new frontiers in scientific research and technological advancements [1]. Quantum machine learning (QML) is one of the central research fields in quantum computing, allowing us to solve various tasks such as classification, regression, and clustering by discovering relationships and patterns between data using quantum computers [2–5]. Recent studies have demonstrated quantum speedups in QML beyond classical machine learning for specific artificially engineered tasks, suggesting the potential of QML [6–8]. While many studies have been devoted to achieving quantum advantages of QML for practical problems, the path to achieving this goal remains unclear [9].

The quantum neural network (QNN) is a promising QML model that combines the principles of quantum information processing and artificial neural networks to enhance the capabilities of data-driven technologies [10–14]. The QNN is represented by a parametrized quantum circuit, which is optimized via training data to solve a given task [15]. Since the efficient simulation of quantum circuits is generally impossi-

ble with classical computers, the QNN can learn the complex features of data that are classically intractable [16,17]. Among various QNN architectures, the quantum convolutional neural network (QCNN) is a leading one that enables classification tasks [18,19] [Fig. 1(a)]. For instance, the QCNN can classify the phases of matter in quantum many-body systems, an important research object in the broad field of physics [20–22]. Due to its high trainability and feasibility [23], the QCNN is particularly suitable for noisy intermediate-scale quantum (NISQ) devices with a limited number of possible gate operations [24].

It is conjectured that achieving quantum advantages in QML requires encoding the prior knowledge of a problem, or inductive bias, into the learning models [25]. In QNNs, the architecture design tailored to prior knowledge is considered crucial to take full advantage of its capabilities [26]. Geometric quantum machine learning (GQML) based on equivariant QNNs, where the symmetry of a problem is encoded in a variational unitary circuit, is one of such prior knowledge-tailored learning models, reducing the parameter space to be searched and enhancing trainability and generalization [27–35]. For example, it was theoretically proved that permutation-equivariant QNNs [36] do not suffer from barren plateaus [37–40], the exponential vanishing of the gradient in a loss function, due to the exponential reduction of parameters to reach overparametrization [41]. This technique would be a powerful tool to achieve quantum advantages with QNNs.

The high resource requirement of the measurement process remains a practical barrier for QNNs to learn data on real quantum computers [15]. During the learning process of QNNs, a predefined loss function is minimized by adjusting

*chinzei.koki@fujitsu.com

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

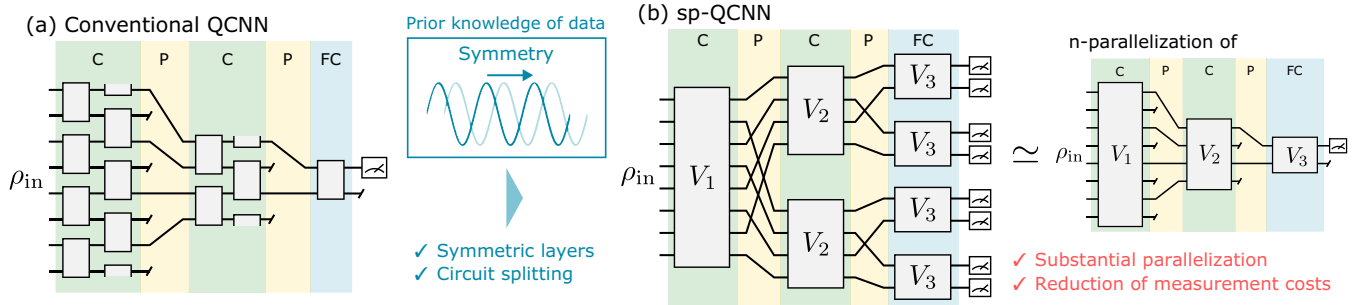


FIG. 1. Basic structures of (a) conventional and (b) sp-QCNNs. C, P, and FC represent convolutional, pooling, and fully connected layers, respectively. (a) In the conventional QCNN, some qubits are discarded at each pooling layer, and only one of the remaining qubits is measured in the end to classify the quantum data. (b) In the sp-QCNN, the translational symmetry of data is used as prior knowledge to design an efficient QML model. The circuit of the sp-QCNN (the left circuit) consists of translationally symmetric layers and splitting structures, allowing us to substantially parallelize the nonsplitting QCNN (the right circuit) to improve the measurement efficiency.

the variational parameters of the circuit. This loss function is computed from a training dataset by measuring specific observables in the parametrized quantum circuit. Therefore, the measurement cost scales with the number of parameters to be optimized and the amount of data to be processed [42]. This situation presents a significant bottleneck when considering large-scale QML applications and the potential of practical quantum advantages [9,43]. To mitigate this measurement requirement, one possible solution is the multiprogramming of quantum computation, which allows multiple circuits to be executed in parallel on different regions of a quantum processor [44–47]. Although this parallelization reduces the total runtime, it increases the required qubit resources, which are limited in current devices.

We address this issue by proposing a QNN architecture called split-parallelizing QCNN (sp-QCNN). This architecture is inspired by GQML and targets translationally symmetric data, such as solid state materials in condensed matter physics. This model exploits the data symmetry as prior knowledge to substantially parallelize the QCNN without increasing the number of qubits, improving the measurement efficiency [Fig. 1(b)]. The circuit of the sp-QCNN consists of two elements: translationally symmetric layers and circuit splitting. First, we impose translational symmetry on the convolutional and fully connected layers to preserve the symmetry of the input state. Second, we split the circuit (rather than discarding some qubits) at the pooling layers and then perform the same unitary operations on each branch in parallel. The combination of this circuit structure and data symmetry substantially parallelizes the conventional QCNN consisting of the same unitary layers, improving the measurement efficiency of local observables and their gradients by a factor of $\mathcal{O}(n)$ (n is the number of qubits throughout this paper).

For verification, we apply the sp-QCNN to a quantum phase recognition task. The results show that the sp-QCNN can improve the measurement efficiency by a factor of $\mathcal{O}(n)$ while achieving sufficient classification performance to recognize the symmetry-protected topological (SPT) phase [48–51]. In training with limited measurement resources, the sp-QCNN with high measurement efficiency can suppress statistical errors in estimating the gradient of the loss function to accelerate the learning process compared to the conven-

tional QCNN. Our model opens up a research direction in the QNN architecture design, contributing to practical quantum advantages in near-term quantum devices that lack sufficient computational resources.

The remainder of this paper is organized as follows. First, Sec. II briefly reviews the QCNN and discusses its computational cost. Section III introduces two key components of the sp-QCNN, translationally symmetric layers and circuit splitting, and clarifies the similarities and differences between the sp-QCNN and the GQML. Section IV shows the advantage of the sp-QCNN, i.e., the improvement of measurement efficiency for local observables and their gradients, based on symmetry. For verification, Sec. V presents the application of the sp-QCNN to a quantum phase recognition task, showing that it can solve the task with sufficient accuracy and improve the measurement efficiency by a factor of $\mathcal{O}(n)$. Finally, Sec. VI summarizes this paper and discusses potential future research directions.

II. REVIEW OF QCNN

The convolutional neural network (CNN) is a celebrated classical machine learning model that solves various tasks, such as image classification [52–54]. The CNN consists of three different types of layers: convolutional, pooling, and fully connected layers. The convolutional layer filters the input data to extract its local features, and the pooling layer coarse-grains the data to leave only relevant information. After the convolutional and pooling layers are alternately applied, the fully connected transformation is applied to the remaining data to produce a final output. For example, in classification problems, the output indicates which class the input data belongs to, and the CNN is trained to correctly classify training data.

The QCNN is a CNN-inspired QNN model that can process quantum data whose dimension is exponentially larger than classical ones and is expected to achieve practical quantum advantages [18,19]. Similar to the CNN, the QCNN consists of convolutional, pooling, and fully connected layers [Fig. 1(a)]. The convolutional layers apply local unitary gates to extract the local features of the input data, and the pooling layers discard some qubits to coarse-grain the quantum information. After alternately applying the two types of layers, we

perform the fully connected unitary, measure the remaining qubits, and obtain an output indicating the data class. In the QCNN, the quantum circuit is characterized by variational parameters, which are optimized to correctly classify training data. Such a variational algorithm is central in the NISQ era, as it works even in a relatively shallow circuit [24].

The QCNN is promising for quantum advantages in NISQ devices because of its two significant features. One is its high feasibility. Since the number of qubits in the QCNN decreases exponentially in each pooling layer, the circuit depth is $\mathcal{O}(\log n)$. This logarithmic depth is advantageous for NISQ devices where the number of possible gate operations is limited. The other feature of the QCNN is its high trainability. In many variational quantum algorithms, the exponential vanishing of the gradient in a loss function, known as the barren plateau phenomenon, prevents scalable optimization [37–40]. Meanwhile, Ref. [23] proved that the QCNN does not suffer from barren plateaus due to the logarithmic depth and the locality of unitary operations and observables. The absence of barren plateaus leads to the high trainability of the QCNN, which is crucial for achieving quantum advantages in QML tasks.

However, the high resource requirement of measurements for optimization presents practical difficulties in QNNs, including QCNN [15]. Let us estimate the required measurement cost in the QCNN. First, we suppose that half of the qubits are discarded at each pooling layer and the number of variational parameters is $\mathcal{O}(n) + \mathcal{O}(n/2) + \mathcal{O}(n/4) + \dots \sim \mathcal{O}(n)$ in total [in common QCNNs, gates acting in parallel share the same parameters, thus the number of independent parameters is $\mathcal{O}(\log n)$, but measuring the gradient of the loss function requires $\mathcal{O}(n)$ cost (see Sec. IV B for details)]. We also let N_{train} , N_{epoch} , and N_{shot} denote the number of training data, maximum epoch (one epoch refers to a complete iteration through a dataset), and the number of measurement shots used per observable, respectively. Then, the total required number of shots during training is $\mathcal{O}(nN_{\text{train}}N_{\text{epoch}}N_{\text{shot}})$. In terms of practicality, the QCNN is not easy to implement for large-scale problems requiring many qubits and a large dataset. Below, we present an architecture of the QCNN that can ideally reduce the required number of shots by a factor of $\mathcal{O}(1/n)$, bringing the QCNN closer to realization.

III. SPLIT-PARALLELIZING QCNN

In this section, we describe the two key components of the sp-QCNN, translationally symmetric layers and circuit splitting, and discuss the relationship between the sp-QCNN and the GQML through symmetry. For simplicity, this work focuses on the case where quantum data is defined on qubits aligned on a one-dimensional lattice. The generalization to arbitrary dimensional lattices is straightforward.

A. Translational symmetry

In the sp-QCNN, we exploit data symmetry as prior knowledge to design an efficient QML model. The target of the sp-QCNN is translationally symmetric data, which is represented by a density matrix ρ_i with the following property:

$$T\rho_iT^\dagger = \rho_i, \quad (1)$$

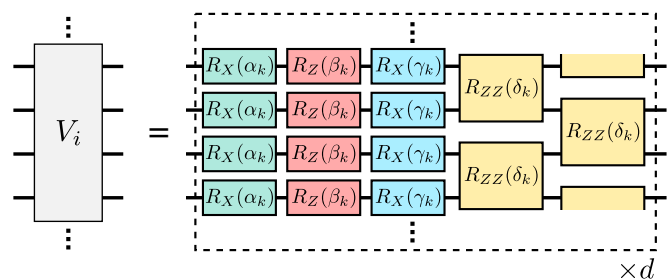


FIG. 2. Example of translationally symmetric unitary layer. Single-qubit rotations are applied in parallel, followed by ZZ rotations on the nearest neighboring qubits. These procedures are repeated d times. The rotation angles are translationally symmetric, and thus the number of independent parameters is $4d$.

where T is the translation operator by one qubit (e.g., $T|100\dots\rangle = |010\dots\rangle$). The most relevant field for the application of sp-QCNN is condensed matter physics, in which translationally symmetric materials such as solids are the largest research topic [55–57]. In Sec. V, we will demonstrate that the sp-QCNN can detect the quantum phases of translationally symmetric many-body states.

To ensure the equivalence of outputs in parallel computation of sp-QCNN, we also impose translational symmetry on each of the convolutional and fully connected layers, whose unitary is denoted by V_i , as follows:

$$TV_iT^\dagger = V_i. \quad (2)$$

As will be shown later, the translational symmetries of the data and the circuit contribute to the substantial parallel computation. We give an example of hardware efficient ansatz respecting the translational symmetry (Fig. 2):

$$V_i = \prod_{k=1}^d R_{ZZ}^{\text{sym}}(\delta_k)R_X^{\text{sym}}(\gamma_k)R_Z^{\text{sym}}(\beta_k)R_X^{\text{sym}}(\alpha_k), \quad (3)$$

where we have defined

$$R_X^{\text{sym}}(\theta) = \prod_{j=1}^n e^{-i\theta X_j}, \quad R_Z^{\text{sym}}(\theta) = \prod_{j=1}^n e^{-i\theta Z_j}, \quad (4)$$

$$R_{ZZ}^{\text{sym}}(\theta) = \prod_{j=1}^n e^{-i\theta Z_j Z_{j+1}}, \quad (5)$$

with the periodic boundary condition $Z_{j+n} = Z_j$. The rotation angles, α_k , β_k , γ_k , and δ_k , are variational parameters to be optimized and do not depend on the qubit position. By construction, V_i is symmetric by one-qubit translation: $[V_i, T] = 0$ (and therefore $[V_i, T^m] = 0$ holds for any integer m). In contrast, the convolutional layer of the conventional QCNN, V_i^{conv} , is symmetric by two or more qubit translations: $[V_i^{\text{conv}}, T] \neq 0$ but, e.g., $[V_i^{\text{conv}}, T^2] = 0$ [Fig. 1(a)].

We note the expressivity and classical simulability of our ansatz. As for the expressivity, the ansatz in Eq. (3) is efficient to implement in quantum hardware but is not general because it cannot express all translationally symmetric unitaries due to the extra inversion symmetry. In principle, the sp-QCNN allows arbitrary translationally symmetric unitaries since we do not impose any constraints on V_i other than translational symmetry (2). For example, unitary operators that break inversion

symmetry, such as $\exp[-i\theta \sum_j X_j Y_{j+1}]$, can be implemented using Trotter decomposition [1]. Then, however, the circuit tends to be deeper and therefore more difficult to implement in near-term devices. Also, extra symmetries of local gates can limit the circuit expressivity [58]. Finding a more expressive and compact ansatz is an important open issue. Meanwhile, it is known that strong symmetry constraints on quantum circuits can lead to classical simulability [59–61]. For example, permutation symmetric quantum circuits are classically simulable in many situations because the dimensions of permutation invariant subspaces are polynomial [62]. Nevertheless, since the translation symmetry is much weaker than the permutation symmetry, we suggest that the translation symmetry does not lead to classical simulability in general. In support of this suggestion, the dimensions of the translationally invariant subspaces are exponentially large. We leave further analysis on it as a future research problem.

B. Circuit splitting

Another key component of the sp-QCNN is circuit splitting. In the conventional QCNN, the pooling layer discards some qubits to coarse-grain the quantum data. In contrast, the sp-QCNN splits the circuit at the pooling layers instead of discarding the qubits, as shown in Fig. 1(b). After splitting, we perform the same operations on each branch and finally measure all the qubits in the computational basis. In some types of quantum computers, such as superconducting [63] and ion-trap devices [64], unitary operations can be performed in parallel, and thus this parallel computation does not significantly increase the runtime.

In this model, we split the circuit such that it is invariant under the translation operation. Figure 3(a) shows an illustrative example of circuit splitting, where the translation operation only swaps the two branches (red and blue lines) but does not modify the overall circuit structure. Here, we give a specific circuit-splitting method for one-dimensional lattice cases. With n as the number of qubits, we choose a prime factor of n , denoted by p , and define $q = n/p$. Then, we introduce splitting in which $n = pq$ qubits are split into p branches [Fig. 3(b)]. First, we divide the qubits into q miniblocks each comprising p qubits in order from top to bottom. Next, we split the circuit such that the j th qubit of the i th miniblock is connected to the i th qubit of the j th branch. By repeating this procedure on each new branch until the number of qubits becomes one, we obtain the entire sp-QCNN circuit. Note that this splitting procedure requires SWAP gates to rearrange the qubits in quantum hardware without all-to-all connectivity (see Appendix A for details).

Due to the translational symmetry of V_i and circuit splitting, the sp-QCNN substantially parallelizes the nonsplitting QCNN that consists of the same V_i [Fig. 1(b)]. For convenience, we define $\langle A \rangle_{\text{ns}}$ and $\langle A \rangle$ as the expectation values of an operator A in the nonsplitting and sp-QCNNs. In the nonsplitting QCNN, we measure one of the remaining qubits in the computational basis and consider its expectation value (i.e., $\langle Z_1 \rangle_{\text{ns}}$) as the output of the QCNN. On the other hand, in the sp-QCNN, we measure all the qubits and regard the average of the n expectation values (i.e., $\langle Z_{\text{avg}} \rangle = \sum_j \langle Z_j \rangle / n$) as the

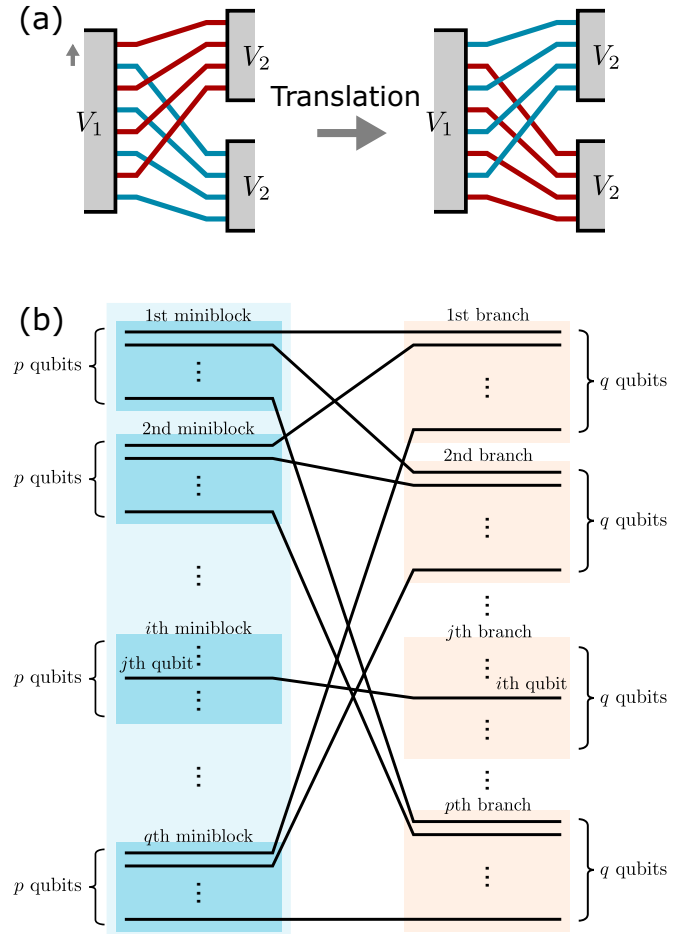


FIG. 3. (a) An illustration of translationally symmetric circuit splitting. In this circuit, the entire circuit structure is invariant under the translation operation. (b) A specific circuit-splitting method. We first divide the qubits into q miniblocks consisting of p qubits and split the circuit such that the j th qubit of the i th miniblock is connected to the i th qubit of the j th branch.

output. In the next section, we will discuss the mechanism and validity of this parallelization in more detail.

C. Relation with geometric quantum machine learning

We consider the sp-QCNN from the viewpoint of GQML or equivariant QNNs [27–36]. The concept of GQML has recently emerged as a potential solution to some critical QML issues associated with trainability and generalization. It leverages the symmetry of a problem as inductive bias and provides a problem-tailored circuit architecture. For example, let us consider the classical task of recognizing whether an image represents a cat. If an image represents a cat, then its rotated image should also represent a cat. In this sense, this task has rotation symmetry. In GQML, such symmetry is encoded in the network architecture. Formally, given a symmetry operation S and an output function $f(\rho)$, the S -invariance of GQML is defined as follows:

$$f(\rho) = f(S\rho S^\dagger) \quad \forall \rho. \quad (6)$$

In other words, the symmetry operation S on the input data never changes the output of GQML. In GQML, the neural network is usually designed based on the equivariant circuit to satisfy this invariance. In theory, GQML significantly enhances the capability of machine learning in several tasks [29,36].

The circuit of the sp-QCNN has the same invariant property as GQML. Let us consider the unitary transformation U of the entire sp-QCNN. Due to the translational symmetry of each V_i and the splitting structure, U itself is translationally symmetric:

$$TUT^\dagger = U. \quad (7)$$

This symmetry leads to the equivariant relation between input and output, $U(T\rho T^\dagger)U^\dagger = T(U\rho U^\dagger)T^\dagger$. That is, the translation operation applied to the input is identical to that applied to the output. We also define $f(\rho) = \text{tr}(U\rho U^\dagger Z_{\text{avg}})$ with an observable $Z_{\text{avg}} = \sum_j Z_j/n$. Then, the equivariant relation and $[Z_{\text{avg}}, T] = 0$ result in

$$f(\rho) = f(T\rho T^\dagger) \quad \forall \rho, \quad (8)$$

which is the T -invariance of GQML (6). Therefore, the sp-QCNN can be seen as applying GQML to the QCNN. This insight suggests that the sp-QCNN can be used to enhance QML capability in tasks where the translation operation on the input data should not change the output. Here, let us clarify the difference between our model and the previously proposed equivariant QCNN [35]. In Ref. [35], the equivariance of the pooling layer in the QCNN is achieved by randomly selecting which qubits to discard based on a given symmetry. The advantage of this conventional model lies in its applicability to various symmetries, while the sp-QCNN is specifically tailored to translational symmetry, imposing equivariance by splitting the circuit. In terms of the number of shots required, our model should be more efficient than the conventional equivariant QCNN by maximizing the utilization of qubit resources.

Our work also provides a different direction for exploiting data symmetry to improve the potential of QML. A critical difference between our problem setting and common ones in GQML is that the input data itself is symmetric in our problem [Eq. (1)], but not in GQML (e.g., the cat image is not rotation invariant). Therefore, each approach brings different benefits. Although the usual GQML improves trainability and generalization, our method reduces measurement costs through substantial parallelization. Thus, the sp-QCNN is particularly advantageous for near-term quantum devices where computational resources are limited.

IV. MEASUREMENT EFFICIENCY IN SP-QCNN

In this section, we describe the parallelization mechanism in the sp-QCNN and show that it can improve the measurement efficiency of local observables and their gradients. We also analytically prove that the improvement factor is $\mathcal{O}(n)$ for a random input state.

A. Measurement efficiency of local observable

First, we show that the translational symmetry of V_i and circuit splitting allow for parallel computation and improve the measurement efficiency of local observables. A key property of the sp-QCNN is the equivalence of expectation values for all the qubits. We recall that the unitary transformation U of the entire sp-QCNN is translationally symmetric [Eq. (7)]. This symmetry leads to

$$\begin{aligned} \langle Z_1 \rangle &= \text{tr}(U\rho U^\dagger Z_1) \\ &= \text{tr}(U(T^\dagger)^{j-1}\rho T^{j-1}U^\dagger Z_1) \\ &= \text{tr}(U\rho U^\dagger Z_j) \\ &= \langle Z_j \rangle, \end{aligned} \quad (9)$$

where ρ is an input state satisfying Eq. (1), and we have used $\rho = (T^\dagger)^{j-1}\rho T^{j-1}$ and $T^{j-1}Z_1(T^\dagger)^{j-1} = Z_j$. This equation indicates the equivalence of the expectation values for all the qubits, i.e., $\langle Z_i \rangle = \langle Z_j \rangle$ for any i and j . This argument can be applied to other single-qubit Pauli operators, leading to $\langle X_i \rangle = \langle X_j \rangle$ and $\langle Y_i \rangle = \langle Y_j \rangle$. Figure 4(a) graphically illustrates this equivalence, which can also be proved by translating the circuit.

This equivalence tells us that the sp-QCNN substantially parallelizes the nonsplitting QCNN that consists of the same V_i , as shown in Fig. 1(b). As mentioned above, we regard the average of the expectation values for all the qubits, $\langle Z_{\text{avg}} \rangle = \sum_j \langle Z_j \rangle / n$, as the output in the sp-QCNN. Meanwhile, we consider the expectation value for only one qubit, $\langle Z_1 \rangle_{\text{ns}}$, as the output in the nonsplitting QCNN. Given the equivalence in Eq. (9), the nonsplitting and sp-QCNNs produce the same results if statistical errors are absent:

$$\langle Z_1 \rangle_{\text{ns}} = \langle Z_{\text{avg}} \rangle. \quad (10)$$

Here, we have used $\langle Z_1 \rangle_{\text{ns}} = \langle Z_1 \rangle$, which can be proved by noticing that the nonsplitting QCNN is a part of the sp-QCNN. In the sp-QCNN, we estimate the output from T measurement shots as follows:

$$\langle Z_{\text{avg}} \rangle_{\text{est}} = \frac{1}{T} \sum_{\ell=1}^T z_{\text{avg}}^{(\ell)} = \frac{1}{nT} \sum_{\ell=1}^T \sum_{j=1}^n z_j^{(\ell)}. \quad (11)$$

Here, $z_j^{(\ell)} = \pm 1$ is the ℓ th measurement outcome at the j th qubit, and we have defined the average of the ℓ th measurement outcomes as $z_{\text{avg}}^{(\ell)} = \sum_j z_j^{(\ell)} / n$. The value of $z_{\text{avg}}^{(\ell)}$ can be a/n ($a \in \{-n, -n+2, \dots, n\}$), corresponding to the measurement outcome of Z_{avg} . We note that the number of outcomes in the sp-QCNN is n times greater than that in the nonsplitting QCNN, where the output is estimated as $\sum_{\ell=1}^T z_1^{(\ell)} / T$. Therefore, the sp-QCNN can reduce the number of shots required to achieve a certain estimation accuracy. Since this argument only relies on the symmetry property of data, the sp-QCNN is general and can be applied to broad tasks with translationally symmetric data.

It is worth noting that the sp-QCNN does not necessarily improve the measurement efficiency by a factor of $\mathcal{O}(n)$. This is because, in each shot, n measurement outcomes are correlated to each other via quantum entanglement. For example, if the output state is the GHZ state $|\psi\rangle = (|000\dots\rangle + |111\dots\rangle) / \sqrt{2}$, then the sp-QCNN

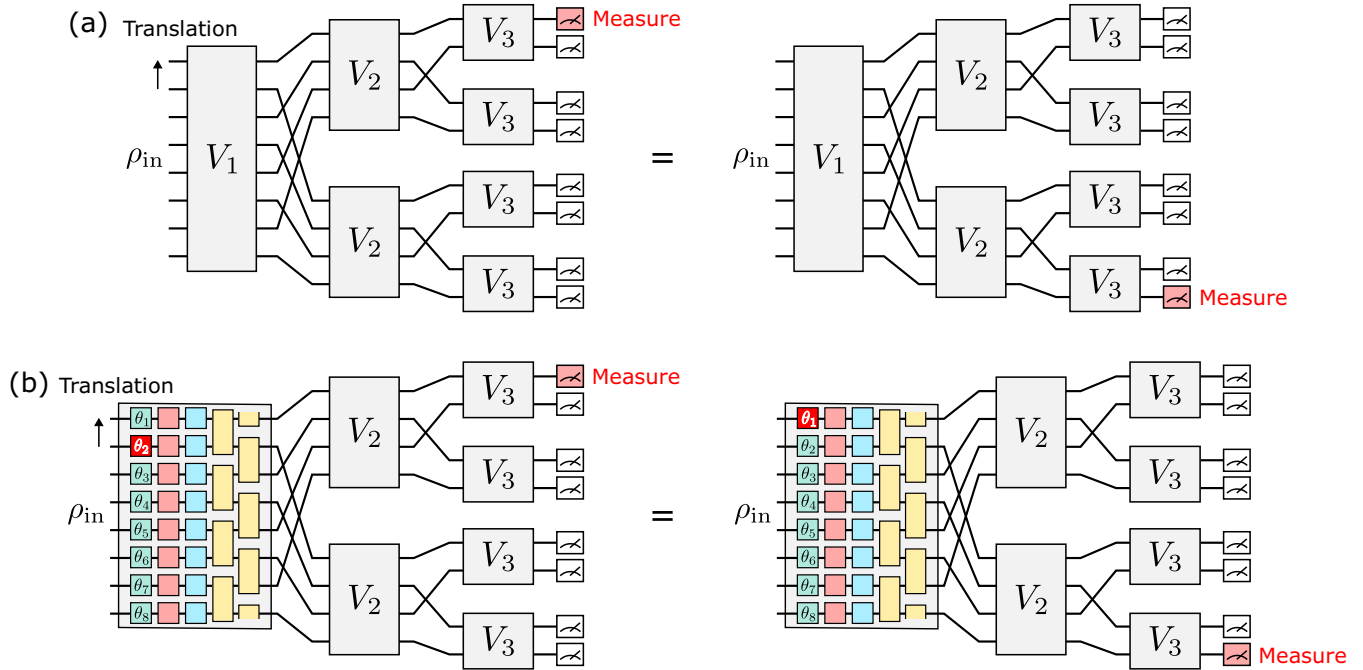


FIG. 4. Mechanism of parallelization in the sp-QCNN. (a) In the sp-QCNN, the expectation value of a local observable is equivalent for all the qubits. This can be proved by virtually translating the entire circuit. The translation does not change the input state and quantum circuit due to their translational symmetry but shifts the position of the measured qubit, showing the equivalence of expectation values at different qubits. (b) The gradient measurement can be parallelized in the sp-QCNN. In accordance with the chain rule, the gradient is the sum of several derivatives, $\partial \langle Z_1 \rangle / \partial \theta = \sum_j \partial \langle Z_1 \rangle / \partial \theta_j$. For example, we suppose that the parameter θ is in the first convolutional layer as shown in the figure (the red boxes denote $\partial / \partial \theta_2$ and $\partial / \partial \theta_1$). Then translating the circuit proves $\partial \langle Z_1 \rangle / \partial \theta_j = \partial \langle Z_{j-2} \rangle / \partial \theta_1$ and thus $\partial \langle Z_1 \rangle / \partial \theta = \sum_j \partial \langle Z_j \rangle / \partial \theta_1$, which can be computed with only two circuits by measuring all the qubits.

does not improve the measurement efficiency at all because the n outcomes are completely correlated and can only provide one bit of information. In contrast, if the output state is the W state $|\psi\rangle = (|100\dots 00\rangle + |010\dots 00\rangle + \dots + |000\dots 01\rangle) / \sqrt{n}$, then the exact expectation value can be obtained with only one shot by measuring all the qubits in the sp-QCNN, whereas many measurements are required in the nonsplitting QCNN. Therefore, how well the sp-QCNN improves the measurement efficiency depends on the details of the problem, such as input data and circuit parameters. Later, we will analytically prove that the sp-QCNN can improve the measurement efficiency by a factor of $\mathcal{O}(n)$ for a typical random input state.

The advantage of the sp-QCNN is illustrated in Fig. 5(a). In actual experiments, we cannot obtain the exact expectation value because of statistical errors. Therefore, it is usually estimated from the mean value of a finite number of measurement outcomes. In the nonsplitting QCNN, the estimated value is generally drawn from the Gaussian distribution with a variance of $\mathcal{O}(1/N_{\text{shot}})$ in accordance with the central limit theorem. In the sp-QCNN, we obtain n measurement outcomes at once and thus expect that the variance scales as $\mathcal{O}(1/nN_{\text{shot}})$, indicating the $\mathcal{O}(n)$ times improvement of measurement efficiency. We note that the sp-QCNN can improve the measurement efficiency of the conventional QCNN, not necessarily other QNNs. Our model further enhances the feasibility of QCNNs, bringing its practical quantum advantages closer to realization.

To quantify the effectiveness of the sp-QCNN, we introduce the relative measurement efficiency:

$$r \equiv \left(\frac{\sigma_0}{\sigma_{\text{sp}}} \right)^2. \tag{12}$$

Here, σ_0 and σ_{sp} are the standard deviations (i.e., square root of variance) of the Gaussians, followed by estimated expectation values in the nonsplitting and sp-QCNNs with the same number of shots. This quantity means that the shot number required to achieve a certain estimation accuracy using the sp-QCNN is $1/r$ times fewer than that using the nonsplitting QCNN. In the next section, we will demonstrate the efficiency of the sp-QCNN for a concrete task using this quantity.

B. Measurement efficiency of gradient

In general, the most costly part of machine learning is the optimization of neural networks using a training dataset, in which the loss function is minimized by tuning the network parameters. In classical machine learning, gradient-based methods are often used for optimization and work well for large-scale problems. Even in QML, gradient-based optimizers are important and powerful tools. However, many measurements are necessary to estimate the gradient in quantum computing [42]. Our architecture makes such gradient measurements efficient.

We first describe a conventional way of measuring the gradient of $\langle Z_1 \rangle$. In many QCNNs, including the sp-QCNN,

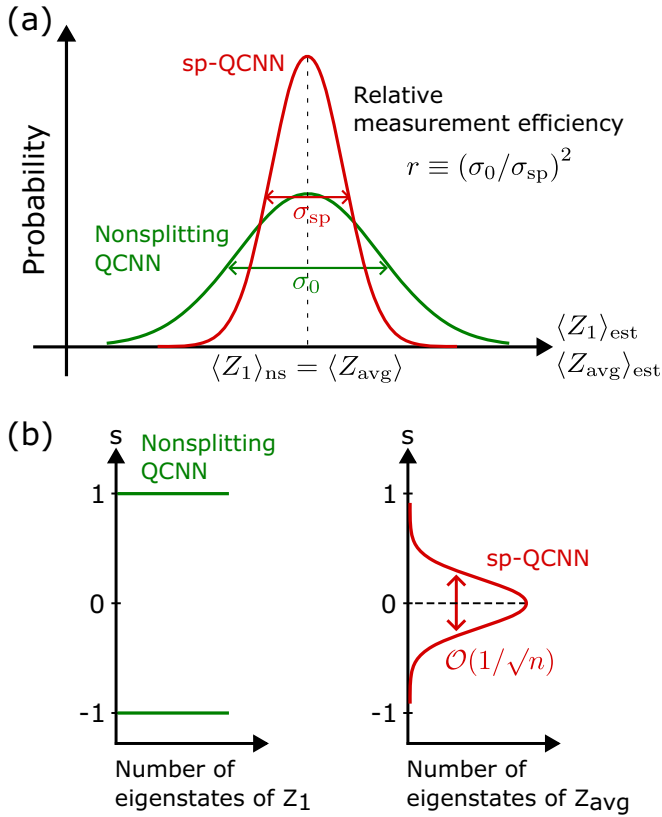


FIG. 5. (a) Quantification of measurement efficiency. In actual experiments, statistical errors arise in estimating the expectation value of an observable. This figure shows the probability distribution of the estimated expectation value. Here, we define the relative measurement efficiency r as the ratio of the variances in the sp-QCNN and the nonsplitting QCNN. (b) Number of eigenstates of Z_1 and Z_{avg} with an eigenvalue s . While the possible measurement outcome is ± 1 in the nonsplitting QCNN (left panel), it is widely distributed in the range of -1 to 1 with a width of $\mathcal{O}(1/\sqrt{n})$ in the sp-QCNN (right panel).

multiple quantum gates share a single variational parameter θ . Here, let m_θ be the number of gates sharing θ in a branch. To calculate the gradient by θ , we suppose that the m_θ gates have different variational parameters from each other, θ_j ($j = 1, \dots, m_\theta$). Thereby, we calculate the gradient with the chain rule as $\partial \langle Z_1 \rangle / \partial \theta = \sum_j (\partial \theta_j / \partial \theta) (\partial \langle Z_1 \rangle / \partial \theta_j) = \sum_j \partial \langle Z_1 \rangle / \partial \theta_j$, where we have used $\partial \theta_j / \partial \theta = 1$. When each gate is parametrized as $e^{-i\theta_j P}$ (P is a Pauli operator), $\partial \langle Z_1 \rangle / \partial \theta_j$ can be measured using the parameter-shift rule, $\partial \langle Z_1 \rangle / \partial \theta_j = \langle Z_1 \rangle_{\theta_j = \theta + \pi/4} - \langle Z_1 \rangle_{\theta_j = \theta - \pi/4}$ [12,65]. Thus, in the sp-QCNN, we can compute the gradient as follows:

$$\frac{\partial \langle Z_1 \rangle}{\partial \theta} = \sum_{j=1}^{m_\theta} \text{tr}(\tilde{U}_{j+} \rho \tilde{U}_{j+}^\dagger Z_1) - \text{tr}(\tilde{U}_{j-} \rho \tilde{U}_{j-}^\dagger Z_1), \quad (13)$$

where $\tilde{U}_{j\pm}$ is the unitary transformation of the sp-QCNN in which $\theta_j = \theta$ is replaced with $\theta_j = \theta \pm \pi/4$. This formula has $2m_\theta$ terms, each of which is usually measured in a different circuit.

The circuit splitting and translational symmetry in the sp-QCNN allow us to compute $\partial \langle Z_1 \rangle / \partial \theta_j$ in parallel, improving

the gradient measurement efficiency. For simplicity, we suppose that each V_i has the form in Eq. (3) and that θ is in the first convolutional layer [Fig. 4(b)]. By translating the entire circuit, we can rewrite each term in Eq. (13) as

$$\begin{aligned} \text{tr}(\tilde{U}_{j\pm} \rho \tilde{U}_{j\pm}^\dagger Z_1) &= \text{tr}(\tilde{U}_{j\pm} T^{j-1} \rho (T^\dagger)^{j-1} \tilde{U}_{j\pm}^\dagger Z_1) \\ &= \text{tr}(\tilde{U}_{1\pm} \rho \tilde{U}_{1\pm}^\dagger Z_{2-j}). \end{aligned} \quad (14)$$

Here, we have used $\rho = T^{j-1} \rho (T^\dagger)^{j-1}$, $(T^\dagger)^{j-1} Z_1 T^{j-1} = Z_{2-j}$, and $(T^\dagger)^{j-1} \tilde{U}_{j\pm} T^{j-1} = \tilde{U}_{1\pm}$. This relation tells us that the derivative of Z_1 by θ_j is identical to that of Z_{2-j} by θ_1 , as illustrated in Fig. 4(b). Thereby, Eq. (13) is reduced to

$$\frac{\partial \langle Z_1 \rangle}{\partial \theta} = \sum_{j=1}^{m_\theta} \text{tr}(\tilde{U}_{1+} \rho \tilde{U}_{1+}^\dagger Z_j) - \text{tr}(\tilde{U}_{1-} \rho \tilde{U}_{1-}^\dagger Z_j), \quad (15)$$

where we have replaced Z_{j-2} with Z_j in the summation. According to this equation, we can obtain the gradient $\partial \langle Z_1 \rangle / \partial \theta$ with just two circuits $\tilde{U}_{1\pm}$ by measuring all the qubits, instead of using $2m_\theta$ circuits that are conventionally necessary.

By generalizing this argument and using the equivalence $\langle Z_{\text{avg}} \rangle = \langle Z_1 \rangle$, we estimate the gradient of the output as follows:

$$\left(\frac{\partial \langle Z_{\text{avg}} \rangle}{\partial \theta} \right)_{\text{est}} = \frac{m_\theta}{nT} \sum_{\ell=1}^T \sum_{j=1}^n [z_{j+}^{(\ell)} - z_{j-}^{(\ell)}], \quad (16)$$

where $z_{j\pm}^{(\ell)}$ is the j th qubit measurement outcome of the ℓ th shot in the parameter-shifted circuit with $\theta_1 = \theta \pm \pi/4$. In our ansatz [Eq. (3)], the factor m_θ/n appears when θ is in the second or later layer. We emphasize that the sp-QCNN enables us to execute n parallel computations even for the gradient estimation, thus accelerating the gradient-based training. Similar to the previous case, the relative measurement efficiency r for the gradient depends on the details of the problem due to the entangled property of the output state.

C. Measurement efficiency for random state

How well the sp-QCNN improves the measurement efficiency depends on the details of the problem. Here, we analytically prove that the efficiency is improved by a factor of $\mathcal{O}(n)$ for a typical state randomly chosen from the T -invariant Hilbert subspace in the limit of $n \rightarrow \infty$.

Let us begin by considering the nonsplitting QCNN, where we measure Z_1 and obtain an outcome $s = \pm 1$ for every measurement. In the limit of $n \rightarrow \infty$, the probability of obtaining an outcome ± 1 is almost $1/2$ for a typical random state because the statistical fluctuations by randomness are negligible due to the exponentially large Hilbert space [this probability distribution is depicted in the left panel of Fig. 5(b)]. Given its Bernoulli distribution, the estimation accuracy of the expectation value is

$$\sigma_0 \sim \mathcal{O}\left(\frac{1}{\sqrt{N_{\text{shot}}}}\right), \quad (17)$$

where N_{shot} is the number of shots.

In the sp-QCNN, we measure all the qubits in the computational basis and regard the mean of the n measurement outcomes as the output of the QCNN [Eq. (11)]. In other

words, we measure $Z_{\text{avg}} = \sum_j Z_j/n$ rather than Z_1 and obtain one of the eigenvalues $s [= \pm 1, \pm(n-2)/n, \dots]$ as an outcome. Also, given that the full unitary transformation U is translationally symmetric, the output state of the sp-QCNN has the same symmetry. The right panel of Fig. 5(b) shows the number of eigenstates of Z_{avg} with an eigenvalue s , $D_n(s)$, on the T -invariant Hilbert subspace. In the limit of $n \rightarrow \infty$, $D_n(s)$ approaches the following asymptotic form (see Appendix B for derivation):

$$D_n(s) \sim \frac{C_n}{(1+s^2)^{n/2}}, \quad (18)$$

where C_n is a constant independent of s . The width of $D_n(s)$ in s is $\mathcal{O}(1/\sqrt{n})$, which finally gives rise to a small estimation error.

Here, we assume that when measuring Z_{avg} for a typical state randomly chosen from the T -invariant subspace, the probability of obtaining an outcome s is proportional to $D_n(s)$. This assumption would be justified in the limit of $n \rightarrow \infty$, where $D_n(s)$ is sufficiently large, and the statistical fluctuations are insignificant. Considering that the width of $D_n(s)$ is $\mathcal{O}(1/\sqrt{n})$, we can estimate the expectation value from N_{shot} experiments with an accuracy

$$\sigma_{\text{sp}} \sim \mathcal{O}\left(\frac{1}{\sqrt{nN_{\text{shot}}}}\right). \quad (19)$$

From the quantification in Eq. (12), the relative measurement efficiency of the sp-QCNN is

$$r = \left(\frac{\sigma_0}{\sigma_{\text{sp}}}\right)^2 \sim \mathcal{O}(n). \quad (20)$$

This result indicates the $\mathcal{O}(1/n)$ times reduction in the number of experiments required to achieve a certain accuracy.

The scaling argument in Eq. (20) is valid in situations where the output state is a random quantum state. Therefore, it may arise in the early stage of the learning process when the parameters of the QCNN are randomly initialized and the output state is approximately random. In Sec. V, we will show that the sp-QCNN exhibits $\mathcal{O}(n)$ scaling for a concrete task in the early stage of learning and, remarkably, even in the final stage.

V. APPLICATION TO QUANTUM PHASE RECOGNITION

In this section, we apply the sp-QCNN to a quantum phase recognition task investigated in Ref. [19] and verify its effectiveness. For the remainder of this paper, we simulate the quantum circuit with Qulacs, an open-source quantum circuit simulator [66].

A. Formulation of problem

Let us consider a one-dimensional cluster Ising model with the periodic boundary condition, whose Hamiltonian is given by

$$H = - \sum_{j=1}^n Z_j X_{j+1} Z_{j+2} - h_1 \sum_{j=1}^n X_j - h_2 \sum_{j=1}^n X_j X_{j+1}, \quad (21)$$

where n is the number of qubits, and X_j, Y_j , and Z_j are the Pauli operators at the j th qubit. This Hamiltonian exhibits SPT [48–51], paramagnetic (PM), and antiferromagnetic (AFM) phases on the h_1 - h_2 plane. The SPT phase is protected by $\mathbb{Z}_2 \times \mathbb{Z}_2$ symmetry characterized by $X_{\text{even(odd)}} = \prod_{j \in \text{even(odd)}} X_j$. The ground state of H , an input state in our task, is translationally symmetric because of $THT^\dagger = H$.

Our task is to recognize the SPT phase using the sp-QCNN. Quantum phase recognition is one of the main applications of the QCNN, and many studies have been conducted with the aim of practical quantum advantages [20–22]. In this task, the sp-QCNN can be applied because the input data (i.e., the ground state of H) is translationally symmetric. For training data, we use 20 ground states of H evenly located on the line of $h_2 = 0$ from $h_1 = 0.05$ to 1.95. Using the Jordan-Wigner transformation [67], we can analytically obtain the exact ground state for $h_2 = 0$, which transits from the SPT to PM phases at $h_1 = 1$. To evaluate the generalization of our method, we test the trained model with 28 data samples, most of which are not included in the training dataset. These samples correspond to ground states at (h_1, h_2) with $h_1 \in \{0.35, 0.65, 0.95, 1.25\}$ and $h_2 \in \{0, \pm 0.5, \pm 1, \pm 1.5\}$. The exact phase diagram on the h_1 - h_2 plane is computed with the density matrix renormalization group (DMRG) [68–71]. In this work, we prepare the input data with exact diagonalization for simplicity. Yet, in actual experiments, other preparation methods must be applied, such as variational quantum eigensolver on a quantum computer and analog-digital transduction from a quantum experiment.

To train our model and evaluate its generalization, we consider the following error as the loss function:

$$\mathcal{L} = \frac{1}{2M} \sum_{i=1}^M (\langle \phi_i | U^\dagger Z_{\text{avg}} U | \phi_i \rangle - y_i)^2, \quad (22)$$

where $|\phi_i\rangle$ and y_i are the training and test data and its corresponding label, and U is the total unitary of the circuit (M is the number of training and test data). Here, we set y_i as 1 if $|\phi_i\rangle$ belongs to the SPT phase and 0 if it does not. We optimize the loss function using the stochastic gradient descent (SGD) method [72]. In SGD, we update the parameters as $\vec{\theta}^{(t+1)} = \vec{\theta}^{(t)} - \eta^{(t)} \nabla \mathcal{L}$, where $\vec{\theta}^{(t)}$ is the parameter vector at optimization step t , and $\nabla \mathcal{L}$ is calculated from only one of the training data at each step. We also decrease the learning rate as $\eta^{(t)} = \eta_0/t$ to stabilize the training and set $\eta_0 = 200$. Besides, to investigate the statistical properties of the sp-QCNN, we simulate the same circuits with N_p different random initial parameter sets. We set $N_p = 50$ in Sec. VB and $N_p = 200$ in Sec. VC.

B. Performance with limited measurement resources

We investigate how the high measurement efficiency of sp-QCNN enhances the machine learning performance in training with limited measurement resources. In such situations, statistical errors in estimating the gradient of the loss function would disturb the training process, reducing accuracy within a limited computational resource. Here, we show that the sp-QCNN can suppress the statistical errors, stabilizing and speeding up the learning process.




FIG. 6. (a and b) Changes in training (top) and test (bottom) loss functions for (a) $n = 8$ and (b) $n = 16$. The orange (blue) solid and dashed lines denote the loss functions with and without statistical errors in the sp-QCNN (conventional QCNN), respectively. The shaded areas are the 10th–90th percentiles of the loss function for 50 sets of random initial parameters at each epoch. We match the number of shots per parameter to obtain the gradient in both QCNNs. (c) Phase diagram predicted by the trained sp-QCNN with statistical errors for $n = 16$ qubits. The color denotes the average magnitude of $\langle Z_{\text{avg}} \rangle$ for 50 sets of initial parameters. The gray dots and dashed green lines denote our training data and phase boundaries computed by DMRG, respectively. In these simulations, we set the depth of each layer as $d = 5$.

As the sp-QCNN circuit, we use the ansatz in Eq. (3) with $d = 5$, where the total number of parameters is 60 for $n = 8$ and 80 for $n = 16$. We also compare the sp-QCNN with the conventional QCNN depicted in Fig. 1(a), where the convolutional and fully connected layers consist of two-qubit unitary gates parametrized as $\prod_{j=1}^{15} e^{-i\theta_j P_j}$ ($P_j = IX, IY, \dots, ZZ$). Since the gates acting in parallel share the same parameters, the number of independent parameters in the conventional QCNN is 75 for $n = 8$ and 105 for $n = 16$. When measuring the gradient in the simulation, we match the shot number per parameter in the conventional and sp-QCNNs for each layer. Here, we use $2m_\theta N_{\text{shot}}$ shots for parameter θ in the sp-QCNN and set $N_{\text{shot}} = 5$.

We first show that our sp-QCNN has comparable classification performance to the conventional QCNN. Figures 6(a) and 6(b) illustrate the changes in the training and test loss functions with and without the statistical errors of gradients for the conventional and sp-QCNNs. Both with and without statistical errors, the final values of the training and test losses in the sp-QCNN (orange lines) are equal to or better than those of the conventional one (blue lines). This result indicates that our model exhibits sufficient expressivity and generalization for this problem compared to the conventional QCNN. Furthermore, our model can accurately predict the unknown phase diagram. Figure 6(c) presents the phase diagram predicted by the trained sp-QCNN with statistical errors for $n = 16$ qubits. It closely aligns with the phase boundary computed by DMRG (shown as dashed lines). Notably, the sp-QCNN can identify the SPT-AFM phase transition, even though this phase transition was not part of the training dataset.

The sp-QCNN effectively suppresses statistical errors and accelerates the learning process while maintaining comparable classification performance to the conventional QCNN. In Figs. 6(a) and 6(b), the loss functions converge rapidly for both QCNNs in the absence of statistical errors (dashed lines). However, in the presence of statistical errors (solid lines), the loss convergence becomes significantly slower in

the conventional QCNN, whereas it remains relatively modest in the sp-QCNN. This fast convergence in the sp-QCNN stems from its high measurement efficiency. While significant statistical errors disturb the rapid and stable optimization in the conventional QCNN, the high measurement efficiency in the sp-QCNN suppresses the statistical errors, stabilizing and accelerating the optimization. As shown in Figs. 6(a) and 6(b), this improvement is more prominent for $n = 16$ compared to $n = 8$, due to the $\mathcal{O}(n)$ improvement in the measurement efficiency. The fast convergence of training is highly effective for near-term quantum devices where a long optimization run is impractical due to limited computational resources.

C. Quantification of measurement efficiency

The previous subsection reveals that the sp-QCNN suppresses the statistical errors in estimating the gradient and accelerates the learning process due to its high measurement efficiency. Here, we numerically quantify the measurement efficiency of sp-QCNN, showing that it improves the measurement efficiency by a factor of $\mathcal{O}(n)$ in the quantum phase recognition task. The measurement efficiency is quantified by the ratio of the variances in the nonsplitting and sp-QCNNs, as shown in Fig. 5 and Eq. (12) (we assume that the nonsplitting QCNN consists of the same unitary V_i as the sp-QCNN). We calculate the measurement efficiency for $n = 8, 12, 16$, and 18 to identify the scaling with respect to the number of qubits. In this verification, the changes in the measurement efficiency are computed during the training process for three typical input states: SPT, PM, and AFM states, which are the eigenstates of H for $(h_1, h_2) = (0, 0), (+\infty, 0),$ and $(0, -\infty)$. We also explore the efficiency of measuring the loss gradient for the first parameter. The unitary circuit consists of the translationally symmetric ansatz (3) with $d = 10$ and is split such that the number of qubits in a branch varies as $8 \rightarrow 4 \rightarrow 2 \rightarrow 1$, $12 \rightarrow 6 \rightarrow 3 \rightarrow 1$, $16 \rightarrow 8 \rightarrow 4 \rightarrow 2 \rightarrow 1$, and $18 \rightarrow 9 \rightarrow 3 \rightarrow 1$ for $n = 8, 12, 16,$ and 18, respectively. Also, to distinguish between the effects of statistical errors on the

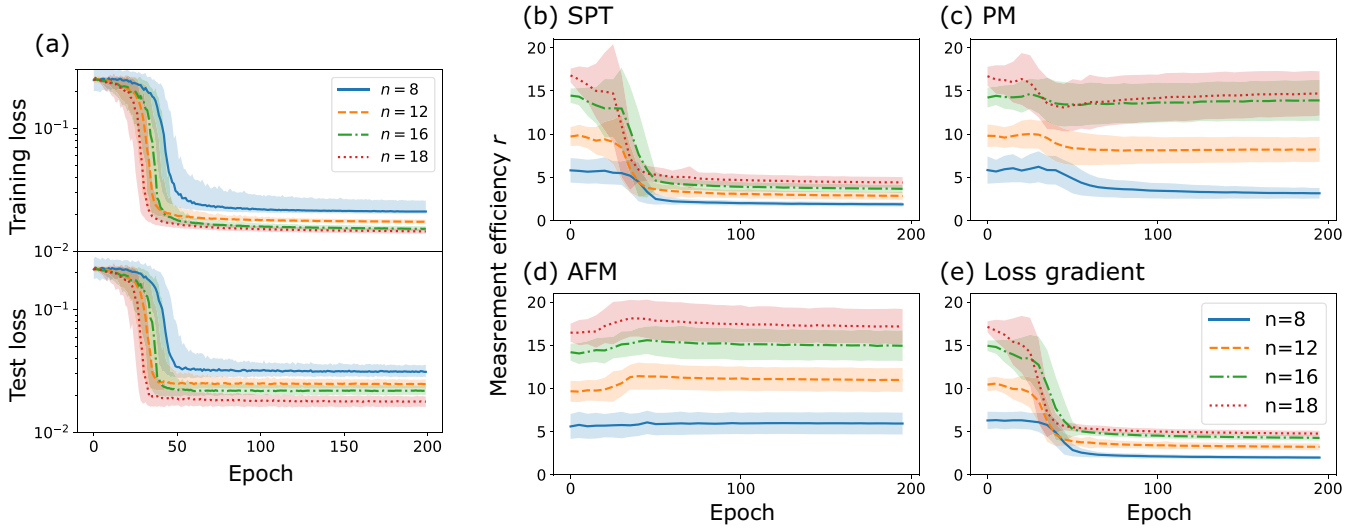


FIG. 7. (a) Changes in training and test loss functions for several numbers of qubits n . The lines and shaded areas depict the median and the 10th–90th percentiles for 200 sets of random initial parameters, respectively. (b–e) Changes in relative measurement efficiency during training. (b)–(d) show the efficiency r for different inputs, SPT, PM, and AFM states, whereas (e) shows the efficiency for measuring the loss gradient by the first parameter. At each epoch, we simulate experiments with 1000 shots 10 000 times, estimate σ_0 and σ_{sp} , and calculate the efficiency $r = (\sigma_0/\sigma_{sp})^2$. The solid lines and shaded areas depict the mean values and standard deviations, respectively, for 200 sets of initial parameters. Except for evaluating the efficiency, we optimize the circuit using the exact expectation value (i.e., without statistical errors).

circuit optimization and the observable measurement, we train the variational circuit with exact gradient (i.e., without statistical errors in estimating gradient) and estimate the measurement efficiency with a finite shot simulation at each epoch. Figure 7(a) shows the training and test losses in this setup, which indicates good trainability and generalization up to $n = 18$.

Figures 7(b)–7(e) show the changes in the relative measurement efficiency r during training for the three inputs and loss gradient. For the PM and AFM states [Figs. 7(c) and 7(d)], the efficiency r is high at the beginning and does not significantly decrease during training. For the SPT state and loss gradient [Figs. 7(b) and 7(e)], r is initially high but decreases as training, finally converging to a small value ($r = 2 - 5$). These results imply that the improvement rate of measurement efficiency strongly depends on the input data, what we measure, and the stage of learning. Even for the

SPT state and loss gradient, the final efficiency is higher than one, indicating that the measurement in the sp-QCNN is more efficient than that in the nonsplitting QCNN.

Figures 8(a) and 8(b) show the relative measurement efficiency with varying the number of qubits n at 0 and 200 epochs for the four cases (cf. Fig. 7). At 0 epoch [Fig. 7(a)], all the data points are nearly aligned on a straight line. This result supports that measurement efficiency is improved by a factor of $\mathcal{O}(n)$ in the early stage of learning and is consistent with the previous argument based on randomness. Even at 200 epoch [Fig. 7(b)], we can fit the data points in straight lines within their error bars, suggesting that the efficiency is also improved by a factor of $\mathcal{O}(n)$ in the final stage of learning. In other words, compared with the nonsplitting QCNN, the sp-QCNN can reduce the number of shots required to achieve a certain estimation accuracy of expectation values by a factor of $\mathcal{O}(1/n)$ throughout the learning process.

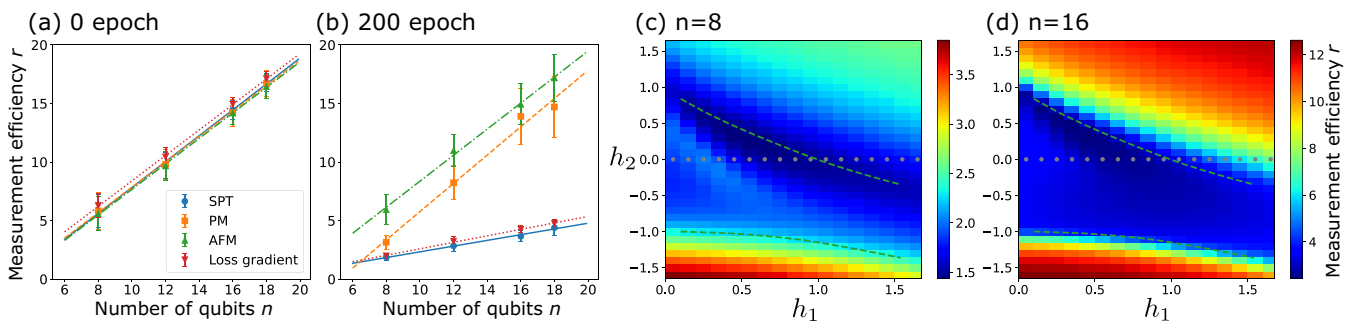


FIG. 8. (a and b) Relative measurement efficiency r with varying the number of qubits n for SPT, PM, AFM, and loss gradient at (a) 0 and (b) 200 epoch in Fig. 7. The four straight lines fit the corresponding types of data points. The error bars denote the standard deviations for 200 sets of initial parameters. (c and d) Relative measurement efficiency r on h_1 - h_2 plane after 200 epochs for (c) $n = 8$ and (d) 16 qubits. The color denotes the magnitude of the relative measurement efficiency. The gray dots and dashed green lines denote our training data and phase boundaries computed by DMRG, respectively.

We also investigate the measurement efficiency for predicting the phase diagram by the trained sp-QCNN with $n = 8$ and 16 qubits [Figs. 8(c) and 8(d)]. By comparing these figures, we notice that the efficiency r for $n = 16$ is more than twice that for $n = 8$ in most areas. This result implies the $\mathcal{O}(n)$ times improvement for prediction. We also observe that the efficiency is low in the SPT phase but relatively high in the PM and AFM phases, a trend evident in Fig. 7 as well. We infer that this phenomenon is due to the following reason. For the SPT state, the expectation value of $\sum_j Z_j/n$ after training is almost one because we have assigned the label as $y_i = 1$ for the SPT phase in the loss function, which means that $U|\phi_{\text{SPT}}\rangle \sim |0\dots 0\rangle$. Given that the sp-QCNN has no advantages for measuring $|0\dots 0\rangle$, the measurement efficiency is not significantly improved. For complete understanding, additional analyses must be conducted in future works.

VI. CONCLUSIONS

In this study, we have proposed a QNN architecture, sp-QCNN, which reduces measurement costs by exploiting the translational symmetry of data as prior knowledge. In the sp-QCNN, we symmetrize and split the QCNN circuit to parallelize the computation, thus improving the measurement efficiency. We have demonstrated the advantage of the sp-QCNN for the quantum phase recognition task: it has high classification performance for this task and can improve the measurement efficiency by a factor of $\mathcal{O}(n)$. In a realistic setting where measurement resources are limited, the sp-QCNN can enhance the speed and stability of the learning process. These results present a possibility for the symmetry-based architecture design of QNN and bring us one step closer to achieving the quantum advantages of the QCNN in near-term quantum devices.

This work offers some research directions for the future. First, finding practical applications of the sp-QCNN is crucial for quantum advantages. A promising candidate is the research of solids, where the sp-QCNN could offer some hints on unsolved problems in condensed matter physics, such as the phase diagrams of the Hubbard model [56] and the kagome antiferromagnetic Heisenberg model [57]. The second direction is further studies of symmetry-based architecture design to reduce measurement costs. Although this work has provided a different approach for QML, its coverage is limited to data with translational symmetry. Hence, generalization to other symmetries, such as the space group, is intriguing and fruitful and may be applied to chemical molecules as well as solid-state materials. The third direction is to find a better ansatz. Although this work establishes the basis of the sp-QCNN, the best V_i for a given problem remains unclear. In general, low expressivity tends to result in poor QML accuracy, while excessively high expressivity can lead to barren plateaus [40,73]. Therefore, finding an ansatz with appropriate expressivity depending on a problem is helpful to realize sp-QCNN experimentally.

Finally, we provide several open issues on the sp-QCNN. This work has shown that the sp-QCNN has sufficient expressivity, trainability, generalization, and measurement efficiency to solve the phase recognition task. However, whether it can solve other complicated tasks remains unclear. In particular,




FIG. 9. An implementation of circuit splitting using local SWAP gates. This figure shows the circuit splitting of n qubits to two branches (n even). The SWAP procedure consists of $(n-2)/2$ steps. For step d , we swap j th and $(j+1)$ th qubits for $j = d+1, d+3, \dots, n-d-1$. This circuit uses $(n^2-2n)/8$ SWAP gates with a depth of $(n-2)/2$.

the translational symmetry of V_i could suppress expressivity and limit solvable tasks in general. Uncovering the possibilities and limitations of the sp-QCNN is an important open issue. For trainability, elucidating whether barren plateaus exist in the sp-QCNN is crucial. In the conventional QCNN, barren plateaus do not appear due to its unique architecture: the logarithmic circuit depth and the locality of unitary operations and observables [23,37–39]. Considering that the sp-QCNN shares these properties with the conventional QCNN, we suggest that no barren plateaus will appear even in the sp-QCNN [74]. The results in this paper show that the training of the sp-QCNN works well up to $n = 18$ qubits, supporting our conjecture. The analysis of measurement efficiency is also an important research issue. Also, while we observed the $\mathcal{O}(n)$ times improvement of measurement efficiency in the phase recognition task of Hamiltonian (21), it remains unclear whether the measurement efficiency is improved by a factor of $\mathcal{O}(n)$ in other problems as well. More thorough analyses are necessary for complete verification.

ACKNOWLEDGMENT

Fruitful discussions with Masatoshi Ishii, Tomochika Kurita, Yuichi Kamata, and Yasuhiro Endo are gratefully acknowledged.

APPENDIX A: IMPLEMENTATION OF CIRCUIT SPLITTING

In real hardware, the implementation of circuit splitting depends on the hardware connectivity. For example, in quantum computers with all-to-all connectivity such as ion trap devices, special processes for circuit splitting are not necessary because the unitary operations on each branch after splitting can be implemented with the qubits being separated. Conversely, in quantum computers where only local entangling operations on neighboring qubits are available, we have to swap the qubits in order to implement the unitary operations after splitting. For concreteness, let us consider a case where n qubits are split into two branches (we assume that n is even). Figure 9

illustrates the SWAP circuit for splitting, which uses $\mathcal{O}(n^2)$ SWAP gates with a depth of $\mathcal{O}(n)$. Although conventional QCNs also require SWAP gates to rearrange the remaining qubits in the pooling layers, our sp-QCNN needs more SWAP gates in general.

APPENDIX B: THE NUMBER OF EIGENSTATES OF Z_{avg}

Here, we derive Eq. (18), where the number of eigenstates of Z_{avg} for translationally symmetric states is

$$D_n(s) \sim \frac{1}{(1+s^2)^{n/2}} \quad (\text{B1})$$

with an eigenvalue s . For convenience, we consider $Z_{\text{tot}} = \sum_j Z_j$ rather than $Z_{\text{avg}} = Z_{\text{tot}}/n$. In the sp-QCNN, we measure Z_{tot} , whose eigenvalues are $\pm n, \pm(n-2), \dots$, and obtain one of the eigenvalues every shot. In addition, the output state is translationally symmetric in the sp-QCNN. Hence, for simplicity, we now focus on the T -invariant eigenspace of Z_{tot} with an eigenvalue z , V_z (i.e., $T|\phi\rangle = |\phi\rangle$ and $Z_{\text{tot}}|\phi\rangle = z|\phi\rangle$ for any $|\phi\rangle \in V_z$). Below, we investigate the dimension of V_z .

To this end, we introduce a cyclic group generated by T ,

$$G_n = \{I, T, T^2, \dots, T^{n-1}\}. \quad (\text{B2})$$

Let M_z be the set of the eigenstates of Z_{tot} with an eigenvalue z in the computational basis (e.g., $M_{n-2} = \{|10\dots 00\rangle, \dots, |00\dots 01\rangle\}$). Then, we define an equivalence relation \sim by G_n in M_z : for $|a\rangle, |b\rangle \in M_z$, $|a\rangle \sim |b\rangle$ holds if and only if $g|a\rangle = |b\rangle$ with $\exists g \in G_n$. We also define the equivalence class of $|a\rangle \in M_z$ as $[a] = \{|x\rangle \in M_z \mid |x\rangle \sim |a\rangle\}$ and the quotient set as $M_z/G_n = \{[a] \mid |a\rangle \in M_z\}$. The elements of M_z/G_n correspond one-to-one to the bases of V_z , such that $|\Psi_i\rangle = \sum_{|\phi\rangle \in [\Phi_i]} |\phi\rangle / \mathcal{N}$, where $|\Psi_i\rangle$ is the base of V_z , $[\Phi_i]$ is the element of M_z/G_n , and \mathcal{N} is the normalization factor ($T|\Psi_i\rangle = |\Psi_i\rangle$ can be easily checked). Therefore, $\dim V_z = |M_z/G_n|$ holds, where $|A|$ is the number of elements in A . Using Burnside's lemma [75], we have $\dim V_z$ as follows:

$$\dim V_z = |M_z/G_n| = \frac{1}{|G_n|} \sum_{g \in G_n} |M_z^g|, \quad (\text{B3})$$

where $M_z^g = \{|\phi\rangle \in M_z \mid g|\phi\rangle = |\phi\rangle\}$. Then, the following theorem holds.

Theorem 1. For $z \neq \pm n$, the following relation holds in the limit of $n \rightarrow \infty$:

$$F_z \equiv \dim V_z \Big/ \frac{1}{n} \binom{n}{\ell_z} \xrightarrow{n \rightarrow \infty} 1 \quad (\text{B4})$$

with $\ell_z = (n+z)/2$. Here, $\binom{n}{\ell_z}$ denotes the binomial coefficient. This theorem states that the asymptotic form of $\dim V_z$ is $\binom{n}{\ell_z}/n$.

Proof. Since $F_z = F_{-z}$ trivially holds, we focus on $-n+2 \leq z \leq 0$, or $1 \leq \ell_z \leq \lfloor n/2 \rfloor$ ($\lfloor \cdot \rfloor$ is the floor function). We first rewrite Eq. (B3) as

$$\begin{aligned} \dim V_z &= \frac{M_z^I}{|G_n|} + \frac{1}{|G_n|} \sum_{g \in G_n \setminus \{I\}} |M_z^g| \\ &= \frac{1}{n} \binom{n}{\ell_z} + \frac{1}{n} \sum_{g \in G_n \setminus \{I\}} |M_z^g|. \end{aligned} \quad (\text{B5})$$




FIG. 10. Illustration for calculating $|M_z^g|$ with $n = 12$, $\ell_z = 4$, and $g = T^3$. Each white (black) circle indicates a single-qubit state of $|0\rangle$ ($|1\rangle$), and $n = 12$ and $\ell_z = 4$ mean that there are eight ($= n - \ell_z$) white and four ($= \ell_z$) black circles in total. Given that the order of g is four ($\because g^4 = I$), we first divide the qubits into four sets, each consisting of three ($= n/\text{ord}(g)$) qubits. In all the sets, the configuration of white and black circles must be the same as each other because of the condition that g does not change the state. Therefore, each set has two white and one ($= \ell_z/\text{ord}(g)$) black circles, and there are three ($= \binom{n/\text{ord}(g)}{\ell_z/\text{ord}(g)}$) possible configurations shown in the figure.

Therefore, F_z is reduced to

$$F_z = 1 + \sum_{g \in G_n \setminus \{I\}} |M_z^g| \Big/ \binom{n}{\ell_z}. \quad (\text{B6})$$

We will evaluate the second term in this equation.

To calculate $|M_z^g|$, we define the order of $g \in G_n$, $\text{ord}(g)$, as the number of elements in the subgroup generated by g (i.e., $\{g^0, g^1, g^2, \dots, g^{k-1}\}$ with $g^k = I$). Note that $\text{ord}(g)$ is a divisor of n . Thereby, $|M_z^g|$ is written as follows:

$$|M_z^g| = \begin{cases} 0 & \ell_z/\text{ord}(g) \notin \mathbb{Z} \\ \binom{n/\text{ord}(g)}{\ell_z/\text{ord}(g)} & \ell_z/\text{ord}(g) \in \mathbb{Z}. \end{cases} \quad (\text{B7})$$

Figure 10 shows a graphical description of Eq. (B7) as an example for $n = 12$, $\ell_z = 4$, and $g = T^3$.

Based on Eq. (B7), one can straightforwardly show that the second term in Eq. (B6) vanishes in the limit of $n \rightarrow \infty$ for $\ell_z = 1, 2$ by noticing that $\text{ord}(g) = 1$ only for $g = I$ and $\text{ord}(g) = 2$ only for $g = T^{n/2}$. Thus, we focus on $3 \leq \ell_z \leq \lfloor n/2 \rfloor$. Because of $\text{ord}(g) \geq 2$ for $g \neq I$, we have

$$|M_z^g| \leq \binom{\lfloor n/2 \rfloor}{\lfloor \ell_z/2 \rfloor}. \quad (\text{B8})$$

This inequality can be shown by considering the properties of binomial coefficients: $\binom{a}{b} < \binom{a'}{b}$ ($a < a'$) and $\binom{a}{b} < \binom{a}{a-b} < \dots < \binom{a}{\lfloor a/2 \rfloor}$ (note that $3 \leq \ell_z \leq \lfloor n/2 \rfloor$). Using Eq. (B8), the

second term in Eq. (B6) is bounded as follows:

$$0 \leq \sum_{g \in G_n \setminus \{I\}} |M_z^g| \left/ \binom{n}{\ell_z} \right. \leq n \binom{\lfloor n/2 \rfloor}{\lfloor \ell_z/2 \rfloor} \left/ \binom{n}{\ell_z} \right. \equiv A_{\ell_z}. \quad (\text{B9})$$

The right-hand side of this inequality, A_{ℓ_z} , approaches zero in $n \rightarrow \infty$ for $3 \leq \ell_z \leq \lfloor n/2 \rfloor$, which can be proven by showing $A_3 \xrightarrow{n \rightarrow \infty} 0$ and $0 < A_{\lfloor n/2 \rfloor} < \dots < A_4 < A_3$ by definition of A_{ℓ_z} . Therefore, we have

$$\sum_{g \in G_n \setminus \{I\}} |M_z^g| \left/ \binom{n}{\ell_z} \right. \xrightarrow{n \rightarrow \infty} 0. \quad (\text{B10})$$

As mentioned above, this limit holds true even for $\ell_z = 1, 2$. Substituting Eq. (B10) to Eq. (B6), we obtain

$$F_z \xrightarrow{n \rightarrow \infty} 1 \quad (\text{B11})$$

with $\ell_z \neq 0, n$. ■

This theorem states that $\dim V_z$ asymptotically approaches $\binom{n}{\ell_z}/n$ (except for $\dim V_n = \dim V_{-n} = 1$). Using Stirling's

formula [$n! \sim \sqrt{2\pi n}(n/e)^n$], we have

$$\dim V_z \sim 2^n \sqrt{\frac{2}{\pi n^3}} D_n(s), \quad (\text{B12})$$

where we have defined

$$D_n(s) \equiv [(1+s)^{1+s+\frac{1}{n}}(1-s)^{1-s+\frac{1}{n}}]^{-n/2} \quad (\text{B13})$$

with $s = z/n$. For large n , $D_n(s)$ rapidly decreases to vanish away from the origin. Therefore, we expand the denominator of $D_n(s)$ in s , obtaining

$$D_n(s) = (1 + [1 + \mathcal{O}(1/n)]s^2 + \mathcal{O}(s^4))^{-n/2} \\ \sim \frac{1}{(1 + s^2)^{n/2}}, \quad (\text{B14})$$

for sufficiently small s . The width of $D_n(s)$ is $\mathcal{O}(1/\sqrt{n})$ in $n \rightarrow \infty$, leading to the $\mathcal{O}(n)$ times improvement of measurement efficiency in the sp-QCNN (see Sec. IV C). Finally, we remark that this discussion is approximately valid for large but finite n while this Appendix considers the limit of $n \rightarrow \infty$. In fact, in Sec. V C, we have observed the clear $\mathcal{O}(n)$ scaling for $n = 18$ at the beginning of training, where the output state is almost random.

-
- [1] M. A. Nielsen and I. L. Chuang, *Quantum Computation and Quantum Information* (Cambridge University Press, Cambridge, 2010).
- [2] N. Wiebe, A. Kapoor, and K. M. Svore, Quantum deep learning, [arXiv:1412.3489](https://arxiv.org/abs/1412.3489).
- [3] M. Schuld, I. Sinayskiy, and F. Petruccione, An introduction to quantum machine learning, *Contemp. Phys.* **56**, 172 (2015).
- [4] J. Biamonte, P. Wittek, N. Pancotti, P. Rebentrost, N. Wiebe, and S. Lloyd, Quantum machine learning, *Nature (London)* **549**, 195 (2017).
- [5] V. Dunjko and H. J. Briegel, Machine learning & artificial intelligence in the quantum domain: a review of recent progress, *Rep. Prog. Phys.* **81**, 074001 (2018).
- [6] 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 (London)* **567**, 209 (2019).
- [7] Y. Liu, S. Arunachalam, and K. Temme, A rigorous and robust quantum speed-up in supervised machine learning, *Nat. Phys.* **17**, 1013 (2021).
- [8] 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**, 2631 (2021).
- [9] M. Schuld and N. Killoran, Is quantum advantage the right goal for quantum machine learning? *PRX Quantum* **3**, 030101 (2022).
- [10] E. Farhi and H. Neven, Classification with quantum neural networks on near term processors, [arXiv:1802.06002](https://arxiv.org/abs/1802.06002).
- [11] J.-G. Liu and L. Wang, Differentiable learning of quantum circuit Born machines, *Phys. Rev. A* **98**, 062324 (2018).
- [12] K. Mitarai, M. Negoro, M. Kitagawa, and K. Fujii, Quantum circuit learning, *Phys. Rev. A* **98**, 032309 (2018).
- [13] M. Benedetti, E. Lloyd, S. Sack, and M. Fiorentini, Parameterized quantum circuits as machine learning models, *Quantum Sci. Technol.* **4**, 043001 (2019).
- [14] M. Schuld, A. Bocharov, K. M. Svore, and N. Wiebe, Circuit-centric quantum classifiers, *Phys. Rev. A* **101**, 032308 (2020).
- [15] 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).
- [16] Y. Du, M.-H. Hsieh, T. Liu, and D. Tao, Expressive power of parametrized quantum circuits, *Phys. Rev. Res.* **2**, 033125 (2020).
- [17] A. Abbas, D. Sutter, C. Zoufal, A. Lucchi, A. Figalli, and S. Woerner, The power of quantum neural networks, *Nat. Comput. Sci.* **1**, 403 (2021).
- [18] E. Grant, M. Benedetti, S. Cao, A. Hallam, J. Lockhart, V. Stojevic, A. G. Green, and S. Severini, Hierarchical quantum classifiers, *npj Quantum Inf.* **4**, 65 (2018).
- [19] I. Cong, S. Choi, and M. D. Lukin, Quantum convolutional neural networks, *Nat. Phys.* **15**, 1273 (2019).
- [20] 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. Zanz, 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, *Nat. Commun.* **13**, 4144 (2022).
- [21] 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).
- [22] 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).
- [23] 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).
- [24] J. Preskill, Quantum computing in the NISQ era and beyond, *Quantum* **2**, 79 (2018).
- [25] J. M. Kübler, S. Buchholz, and B. Schölkopf, The inductive bias of quantum kernels, [arXiv:2106.03747](https://arxiv.org/abs/2106.03747).
- [26] 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).
- [27] M. M. Bronstein, J. Bruna, T. Cohen, and P. Veličković, Geometric deep learning: Grids, groups, graphs, geodesics, and gauges, [arXiv:2104.13478](https://arxiv.org/abs/2104.13478).
- [28] G. Verdon, T. McCourt, E. Luzhnica, V. Singh, S. Leichenauer, and J. Hidary, Quantum graph neural networks, [arXiv:1909.12264](https://arxiv.org/abs/1909.12264).
- [29] 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).
- [30] M. Larocca, F. Sauvage, F. M. Sbahti, G. Verdon, P. J. Coles, and M. Cerezo, Group-invariant quantum machine learning, *PRX Quantum* **3**, 030341 (2022).
- [31] 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).
- [32] A. Skolik, M. Cattelan, S. Yarkoni, T. Bäck, and V. Dunjko, Equivariant quantum circuits for learning on weighted graphs, [arXiv:2205.06109](https://arxiv.org/abs/2205.06109).
- [33] F. Sauvage, M. Larocca, P. J. Coles, and M. Cerezo, Building spatial symmetries into parameterized quantum circuits for faster training, *Quantum Sci. Technol.* **9**, 015029 (2024).
- [34] 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, [arXiv:2210.07980](https://arxiv.org/abs/2210.07980).
- [35] Q. T. Nguyen, L. Schatzki, P. Braccia, M. Ragone, P. J. Coles, F. Sauvage, M. Larocca, and M. Cerezo, Theory for equivariant quantum neural networks, [arXiv:2210.08566](https://arxiv.org/abs/2210.08566).
- [36] L. Schatzki, M. Larocca, F. Sauvage, and M. Cerezo, Theoretical guarantees for permutation-equivariant quantum neural networks, *npj Quantum Inf* **10**, 12 (2024).
- [37] J. R. McClean, S. Boixo, V. N. Smelyanskiy, R. Babbush, and H. Neven, Barren plateaus in quantum neural network training landscapes, *Nat. Commun.* **9**, 4812 (2018).
- [38] 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**, 1791 (2021).
- [39] C. O. Marrero, M. Kieferová, and N. Wiebe, Entanglement-induced barren plateaus, *PRX Quantum* **2**, 040316 (2021).
- [40] Z. Holmes, K. Sharma, M. Cerezo, and P. J. Coles, Connecting ansatz expressibility to gradient magnitudes and barren plateaus, *PRX Quantum* **3**, 010313 (2022).
- [41] 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).
- [42] A. Abbas, R. King, H.-Y. Huang, W. J. Huggins, R. Movassagh, D. Gilboa, and J. R. McClean, On quantum backpropagation, information reuse, and cheating measurement collapse, [arXiv:2305.13362](https://arxiv.org/abs/2305.13362).
- [43] H.-Y. Liu, Z.-Y. Chen, T.-P. Sun, C. Xue, Y.-C. Wu, and G.-P. Guo, Can variational quantum algorithms demonstrate quantum advantages? Time really matters, [arXiv:2307.04089](https://arxiv.org/abs/2307.04089).
- [44] P. Das, S. S. Tannu, P. J. Nair, and M. Qureshi, A case for multi-programming quantum computers, in *Proceedings of the 52nd Annual IEEE/ACM International Symposium on Microarchitecture*, MICRO '52 (ACM Press, New York, NY, 2019), pp. 291–303.
- [45] L. Liu and X. Dou, QuCloud: A new qubit mapping mechanism for multi-programming quantum computing in cloud environment, in *2021 IEEE International Symposium on High-Performance Computer Architecture (HPCA)* (IEEE, Seoul, Korea, 2021), pp. 167–178.
- [46] S. Niu and A. Todri-Sanial, How parallel circuit execution can be useful for NISQ computing? [arXiv:2112.00387](https://arxiv.org/abs/2112.00387).
- [47] S. Niu and A. Todri-Sanial, Enabling multi-programming mechanism for quantum computing in the NISQ era, *Quantum* **7**, 925 (2023).
- [48] Z.-C. Gu and X.-G. Wen, Tensor-entanglement-filtering renormalization approach and symmetry-protected topological order, *Phys. Rev. B* **80**, 155131 (2009).
- [49] F. Pollmann, A. M. Turner, E. Berg, and M. Oshikawa, Entanglement spectrum of a topological phase in one dimension, *Phys. Rev. B* **81**, 064439 (2010).
- [50] X. Chen, Z.-C. Gu, and X.-G. Wen, Classification of gapped symmetric phases in one-dimensional spin systems, *Phys. Rev. B* **83**, 035107 (2011).
- [51] F. Pollmann and A. M. Turner, Detection of symmetry-protected topological phases in one dimension, *Phys. Rev. B* **86**, 125441 (2012).
- [52] Y. LeCun, Y. Bengio, and G. Hinton, Deep learning, *Nature (London)* **521**, 436 (2015).
- [53] Y. Lecun and Y. Bengio, Convolutional networks for images, speech, and time-series, in *The Handbook of Brain Theory and Neural Networks*, edited by M. A. Arbib (MIT Press, Cambridge, 1995).
- [54] A. Krizhevsky, I. Sutskever, and G. E. Hinton, ImageNet classification with deep convolutional neural networks, in *Advances in Neural Information Processing Systems* (Curran Associates, Lake Tahoe, Nevada, 2012), Vol. 25.
- [55] B. Bauer, S. Bravyi, M. Motta, and Garnet K.-L. Chan, Quantum algorithms for quantum chemistry and quantum materials science, *Chem. Rev.* **120**, 12685 (2020).
- [56] D. P. Arovas, E. Berg, S. A. Kivelson, and S. Raghu, The Hubbard model, *Annu. Rev. Condens. Matter Phys.* **13**, 239 (2022).
- [57] S. Yan, D. A. Huse, and S. R. White, Spin-liquid ground state of the $S = 1/2$ kagome Heisenberg antiferromagnet, *Science* **332**, 1173 (2011).
- [58] I. Marvian, Restrictions on realizable unitary operations imposed by symmetry and locality, *Nat. Phys.* **18**, 283 (2022).
- [59] I. Kerenidis, J. Landman, and N. Mathur, Classical and quantum algorithms for orthogonal neural networks, [arXiv:2106.07198](https://arxiv.org/abs/2106.07198).
- [60] M. L. Goh, M. Larocca, L. Cincio, M. Cerezo, and F. Sauvage, Lie-algebraic classical simulations for variational quantum computing, [arXiv:2308.01432](https://arxiv.org/abs/2308.01432).
- [61] M. Cerezo, M. Larocca, D. García-Martín, N. L. Diaz, P. Braccia, E. Fontana, M. S. Rudolph, P. Bermejo, A. Ijaz,

- S. Thanasilp, E. R. Anschuetz, and Z. Holmes, Does provable absence of barren plateaus imply classical simulability? Or, why we need to rethink variational quantum computing, [arXiv:2312.09121](https://arxiv.org/abs/2312.09121).
- [62] E. R. Anschuetz, A. Bauer, B. T. Kiani, and S. Lloyd, Efficient classical algorithms for simulating symmetric quantum systems, *Quantum* **7**, 1189 (2023).
- [63] C. Song, K. Xu, W. Liu, C.-P. Yang, S.-B. Zheng, H. Deng, Q. Xie, K. Huang, Q. Guo, L. Zhang, P. Zhang, D. Xu, D. Zheng, X. Zhu, H. Wang, Y.-A. Chen, C.-Y. Lu, S. Han, and J.-W. Pan, 10-qubit entanglement and parallel logic operations with a superconducting circuit, *Phys. Rev. Lett.* **119**, 180511 (2017).
- [64] C. Figgatt, A. Ostrander, N. M. Linke, K. A. Landsman, D. Zhu, D. Maslov, and C. Monroe, Parallel entangling operations on a universal ion-trap quantum computer, *Nature (London)* **572**, 368 (2019).
- [65] M. Schuld, V. Bergholm, C. Gogolin, J. Izaac, and N. Killoran, Evaluating analytic gradients on quantum hardware, *Phys. Rev. A* **99**, 032331 (2019).
- [66] Y. Suzuki, Y. Kawase, Y. Masumura, Y. Hiraga, M. Nakadai, J. Chen, K. M. Nakanishi, K. Mitarai, R. Imai, S. Tamiya, T. Yamamoto, T. Yan, T. Kawakubo, Y. O. Nakagawa, Y. Ibe, Y. Zhang, H. Yamashita, H. Yoshimura, A. Hayashi, and K. Fujii, Qulacs: A fast and versatile quantum circuit simulator for research purpose, *Quantum* **5**, 559 (2021).
- [67] P. Jordan and E. Wigner, Über das paulische äquivalenzverbot, *Z. Phys.* **47**, 631 (1928).
- [68] S. R. White, Density matrix formulation for quantum renormalization groups, *Phys. Rev. Lett.* **69**, 2863 (1992).
- [69] S. R. White, Density-matrix algorithms for quantum renormalization groups, *Phys. Rev. B* **48**, 10345 (1993).
- [70] U. Schollwöck, The density-matrix renormalization group in the age of matrix product states, *Ann. Phys.* **326**, 96 (2011).
- [71] M. Fishman, S. White, and E. Stoudenmire, The ITensor software library for tensor network calculations, *SciPost Phys. Codebases* **4** (2022).
- [72] H. Robbins and S. Monro, A stochastic approximation method, *Ann. Math. Statist.* **22**, 400 (1951).
- [73] E. R. Anschuetz and B. T. Kiani, Quantum variational algorithms are swamped with traps, *Nat. Commun.* **13**, 7760 (2022).
- [74] C. Tüysüz, G. Clemente, A. Crippa, T. Hartung, S. Kühn, and K. Jansen, Classical splitting of parametrized quantum circuits, *Quantum Mach. Intell.* **5**, 34 (2023).
- [75] J. J. Rotman, *An Introduction to the Theory of Groups* (Springer, New York, 1994).