



## OPEN Deep quanvolutional neural networks with enhanced trainability and gradient propagation

Muhammad Kashif<sup>1,2✉</sup> & Muhammad Shafique<sup>1,2</sup>

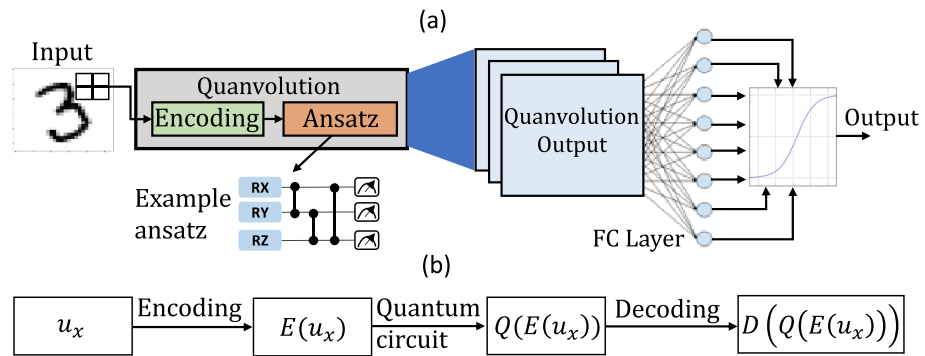
In this paper, we explore methods to enhance the performance of one of the frequently used variants of Quantum Convolutional Neural Networks, known as Quanvolutional Neural Networks (QuNNs) by introducing trainable quanvolutional layers and addressing the challenges associated with training multi-layered or deep QuNNs. Traditional QuNNs mostly rely on static (non-trainable) quanvolutional layers, limiting their feature extraction capabilities. Our approach enables the training of these layers, significantly improving the scalability and learning potential of QuNNs. However, multi-layered deep QuNNs face difficulties in gradient-based optimization due to limited gradient flow across all the layers of the network. To overcome this, we propose Residual Quanvolutional Neural Networks (ResQuNNs), which utilize residual learning by adding skip connections between quanvolutional layers. These residual blocks enhance gradient flow throughout the network, facilitating effective training in deep QuNNs, thus *enabling deep learning in QuNNs*. Moreover, we provide empirical evidence on the optimal placement of these residual blocks, demonstrating how strategic configurations improve gradient flow and lead to more efficient training. Our findings represent a significant advancement in quantum deep learning, opening new possibilities for both theoretical exploration and practical quantum computing applications.

**Keywords** Quantum neural networks, Residual learning, Trainability, Deep learning

Quantum Machine Learning (QML) combines quantum computing with machine learning to process and analyze data using quantum algorithms and circuits<sup>1</sup>. By leveraging quantum phenomena such as superposition and entanglement, QML aims to provide computational advantages over classical methods, especially for complex, high-dimensional problems<sup>2–4</sup>. In QML, Quantum Convolutional Neural Networks (QCNNs) represent an innovative fusion of quantum computing principles with traditional convolutional neural network (CNN) architectures<sup>5,6</sup>. These hybrid classical-quantum models utilize parameterized quantum circuits (PQCs) *a.k.a* ansatz, composed of parameterized rotation gates such as  $R_x$ ,  $R_y$ , or  $R_z$ , and entangling gates such as  $CX$  and  $CZ$ , functioning analogously to convolutional and pooling operations in classical CNNs<sup>7,8</sup>. The optimization, however, is performed on classical computers<sup>9–12</sup>. A notable advantage of QCNNs lies in their ability to harness the vast Hilbert space offered by quantum mechanics, which surpasses the capabilities of classical CNNs<sup>13</sup>. This feature allows QCNNs to capture spatial relationships in image data more effectively. Furthermore, the potential future inclusion of higher-level quantum systems, such as qutrits or ququads<sup>14,15</sup>, promises even more sophisticated image comprehension capabilities.

There are two primary variants of QCNNs, each leveraging quantum mechanics principles like superposition, entanglement, and interference. The first type, Quantum-Inspired CNNs, proposed in<sup>7</sup>, mirrors the structure of classical CNNs but replaces convolutional and pooling layers with deeper PQCs. This design demands higher qubit coherence to fully exploit quantum parallelism and entanglement, posing a challenge for current quantum hardware. Nevertheless, a significant research has been dedicated to explore this very architectural framework for various applications<sup>9,16</sup>. The second type, Quanvolutional Neural Networks (QuNNs), first introduced in<sup>17</sup>, and visually depicted in Fig. 1a, focuses on quantum-enhanced convolutional layers. These networks encode classical image pixels into quantum states, apply unitary transformations via PQCs, and then measure the qubits

<sup>1</sup>eBrain Lab, Division of Engineering, New York University Abu Dhabi, PO Box 129188, Abu Dhabi, United Arab Emirates. <sup>2</sup>Center for Quantum and Topological Systems, NYUAD Research Institute, New York University, Abu Dhabi, United Arab Emirates. ✉email: muhammadkashif@nyu.edu



**Fig. 1.** (a) A typical illustration of quanvolutional neural network architecture, (b) simplified workflow of quanvolutional layer in QuNNs.

Ref	Trainable quanvolutional layers	Multiple quanvolutional layers
22	×	×
23	×	×
24	×	×
25	×	×
26	×	×
27	×	×
28	✓	×
29	✓	×
<b>Ours</b>	✓	✓

**Table 1.** Recent state-of-the-art QuNNs employing trainable and untrainable quanvolution layers.


to create output image channels, which are subsequently processed by classical fully connected layers. This approach offers a more feasible implementation with current quantum technology. Hence, we also delve into QuNNs with a focus on identifying the challenges faced in training QuNNs and proposing effective solutions to address these issues.

An abstract view of the workflow of quanvolution layer in QuNNs is presented in Fig. 1b, which is four step process: (1) the original image  $U$  is divided into spatially located subsections  $u_x$ . Each subsection is a  $d \times d$  matrix with  $d > 1$ , (2) these image subsections are then encoded into a quantum state ( $q$ ), e.g., for an encoding function  $E$ , the quantum state can be defined as  $q_x = E(u_x)$ . We encode the input features into the qubit rotation angles using  $R_y(\theta)$ . (3) the encoded state  $q_x$  then undergoes a unitary transformation via a quantum circuit *a.k.a* ansatz  $Q$  which can be described as;  $f_x = Q(q_x) = Q(E(u_x))$ , (4) decoding is then performed via measurements which can be described as,  $D(f_x) = D(Q(E(u_x)))$ . By repeating these steps for different subsections, the full input image is traversed, producing an outcome organized as a multi-channel image. The output of this quanvolution is usually processed by a classical neural network to produce the final output.

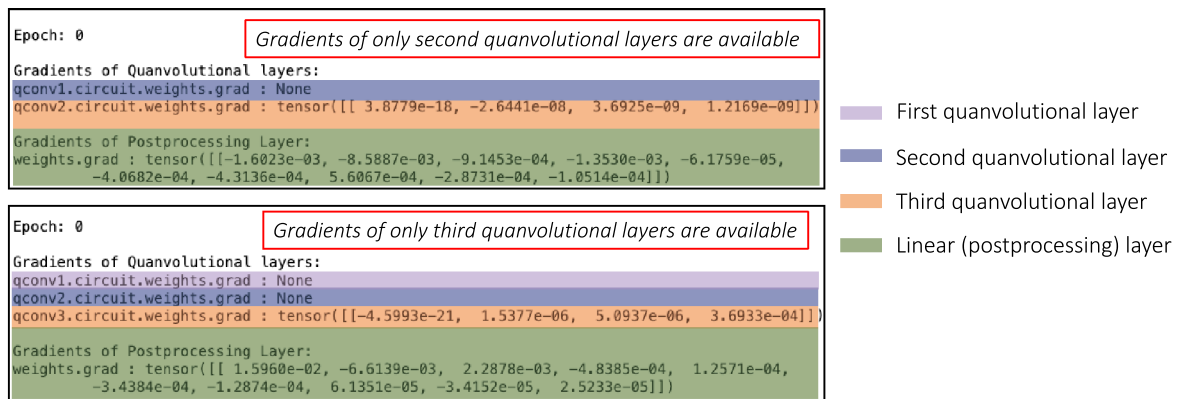
The QuNN architecture depicted in Fig. 1, widely adopted in state-of-the-art research for various applications such as 2D and 3D radiological image classification<sup>18</sup>, high energy physics data analysis<sup>19</sup>, and object detection and classification<sup>20,21</sup>. In these contexts, the quantum circuit in QuNNs primarily functions as a feature extractor and is not trained during the quanvolution operation. Consequently, the learning process is primarily driven by the subsequent classical network. Although there are recent studies that employ trainable quantum circuits in quanvolutional layers, to the best of our knowledge, no existing literature addresses multi-layered QuNNs. Table 1 some summarizes recent works utilizing QuNNs with both trainable and non-trainable quanvolutional layers. To achieve full scalability in QuNNs, it is crucial to ensure the trainability of underlying quanvolutional layers in multi-layered architectures.

### Motivational case study

In Fig. 2, we assess the impact of trainable versus untrainable quanvolutional layers in QuNNs. To this end, we train QuNNs with both trainable (allowing weight updates during backpropagation) and untrainable (frozen weights, no updates) quanvolutional layers. This training is performed on 1000 randomly chosen images from MNIST dataset over 30 iterations, using an 80% – 20% training-validation split. The Adam optimizer with a 0.01 learning rate was employed. We utilized a  $2 \times 2$  kernel resulting in 4-qubit quanvolutional layers, with each layer comprising of 4 parameterized  $R_y$  gates and 3 *CNOT* gates for (nearest neighbor) qubit entanglement. The overall depth of quantum circuit in the quanvolutional layer is set to 2, accounting to a total of 8 single and 6 two-qubit gates.



**Fig. 2.** Motivational analysis showing the impact *Trainable* quantum convolutional layers. When the quantum convolutional layer is untrainable (UT) and only the classical linear layer is trained (T), the QuNN yields suboptimal performance. When the quantum convolutional layer is made trainable, it significantly enhanced the training performance. *Quantum* quantum convolutional layer, *Linear* fully connected layer at the end.



**Fig. 3.** Gradients Propagation in multi-layered QuNNs with Two (upper panel) and Three (lower panel) Quantum convolutional Layers.

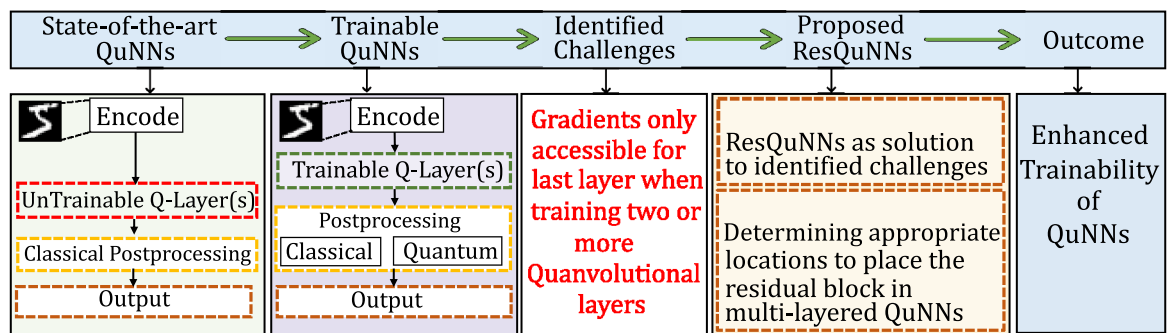
Results showed suboptimal performance when quantum convolutional layers were untrainable and only the final fully connected (FC) layer was trained. Conversely, allowing weight updates in the quantum convolutional layers significantly improved learning, enhancing model accuracy by approximately 36%. *This empirical evidence robustly substantiates the hypothesis that facilitating the trainability of quantum convolutional layers significantly augments the model's overall learning efficacy.*

While a single trainable quantum convolutional layer enhances QuNNs' training efficacy, as depicted in Fig. 2, scalability of QuNNs requires multiple trainable layers. However, challenges emerge when multiple quantum convolutional layers are trained concurrently, as gradients are exclusively accessible only in the last quantum convolutional layer, thus limiting the optimization process only to this layer. The gradient accessibility in multi-layered QuNNs is hindered because each quantum layer performs qubit measurement, thus collapsing the quantum state and breaking the chain of differentiability. Figure 3 demonstrates this by showing gradients for two and three trainable quantum convolutional layers. We observe that the gradients of only the final quantum convolutional layer are accessible, while those of the preceding layers are effectively *None*. The restricted gradient flow means that in multi-layered QuNNs, only the last quantum convolutional layer contributes in the learning process, leaving earlier layers remain uninvolved in optimization. This presents a significant challenge in effectively training multi-layered QuNNs, thereby necessitating innovative approaches to address this challenge for scalable and effective QuNN architectures.

### Our contributions

Our primary contributions encompass the introduction of trainable quantum convolutional layers, identification of challenges with multiple quantum convolutional layers, the development of Residual Quantum convolutional Neural Networks (ResQuNNs) to address these challenges, and empirical findings on the strategic placement of residual blocks to further enhance training performance. An overview of our contributions is shown in Figure 4.

- Trainable quantum convolutional layers. We first make the quantum convolutional layers trainable in QuNNs, addressing the challenge of enhancing their adaptability within QuNNs and improving their training performance.
- Challenges with multiple quantum convolutional layers. We identify a key challenge in optimizing multiple quantum convolutional layers, highlighting the complexities associated with accessing gradients across these layers. In QuNN architectures with multiple trainable quantum convolutional layers, the gradients do not flow through all the layers of the network, making deeper QuNNs unsuitable for gradient-based optimization techniques.



**Fig. 4.** Overview of our contributions.

- Residual quantum convolutional neural networks (ResQuNN). To overcome the gradient accessibility issue, this paper introduces Residual Quantum Convolutional Neural Networks (ResQuNN), drawing inspiration from classical residual neural networks.
- Residual blocks for gradient accessibility. In ResQuNNs, the residual blocks between quanvolutional layers are incorporated to facilitate comprehensive gradient access throughout the network, thereby enabling the deep learning in QuNNs and eventually improving training performance of QuNNs. These residual blocks combine the output of existing quanvolutional layers with their own or previous layer's input before forwarding it to the subsequent quanvolutional layer.
- Strategic placement of residual blocks. We conducted extensive experiments to determine appropriate locations for inserting residual blocks in ResQuNNs which ensure comprehensive gradient throughout the network.
- Classical and quantum postprocessing. State-of-the-art QuNNs utilize classical networks to postprocess the output of quanvolutional layers. However, to demonstrate the effectiveness of the residual approach in multi-layered QuNNs (with no influence on the overall learning from classical layers), other than classical post-processing layer, we employ the quantum circuit itself for postprocessing the results of the quanvolutional layers.

### Related work

Recently, some studies have explored the application of residual connections in quantum neural networks (QNNs). In<sup>30</sup>, residual connections are incorporated into a simple feedforward QNN architecture to address the problem of barren plateaus (BPs). BPs are characterized by the exponential vanishing of gradients in parameterized quantum circuits, which are essential components of QNNs, as the number of qubits increases. By partitioning the conventional QNN architectures into multiple nodes, and then by adding residual connections between these nodes,<sup>30</sup> successfully mitigates the BP issue. Our work diverges from<sup>30</sup> in two significant ways: (1) We employ residual approach in a different architecture, specifically QuNNs, and (2) we address a distinct underlying problem, i.e., restricted gradient flow in QuNNs, which occurs regardless of the number of qubits unlike BPs which occurs with increasing number of qubits. To overcome the partial gradient flow in QuNNs, we apply residual connections.

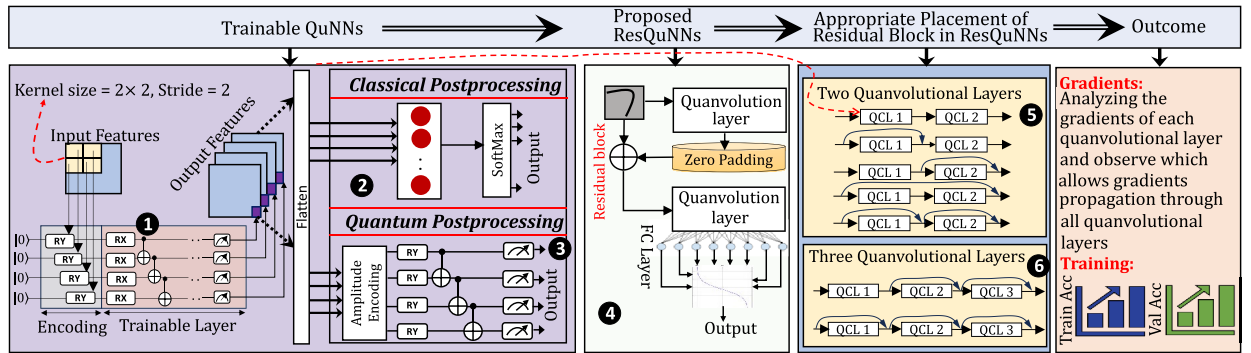
In another recent work<sup>31</sup>, the concept of residual connections is used in QNNs. The proposed quantum residual neural networks (QResNets) leverage auxiliary qubits in data-encoding and trainable blocks to extend frequency generation forms and improve the flexibility of adjusting Fourier coefficients, leading to enhanced spectral richness and expressivity of parameterized quantum circuits at the cost of greater qubit utilization in the form of auxiliary qubits, which is not suitable for NISQ devices. In contrast, our work focuses on QuNNs, addressing the critical challenges of restricted gradient flow in deep, multi-layered QuNNs without using any additional qubits. We introduce trainable quanvolutional layers, significantly increasing the scalability and learning potential of QuNNs. To resolve complexities in gradient-based optimization for QuNNs with multiple layers, we propose ResQuNNs that utilize residual learning and skip connections to facilitate gradient flow across multiple trainable quanvolutional layers. While QResNets enhance frequency generation and flexibility in quantum circuits, our work specifically targets the scalability and learning potential of QuNNs, providing practical solutions for efficient training in deep quantum neural networks.

### ResQuNN methodology

The detailed overview of ResQuNN's methodology is presented in Fig. 5. Below we discuss different steps of our methodology in detail.

#### Trainable quanvolutional layers in QuNNs

The first step is to make the quanvolutional layers trainable. The majority of works in the literature utilizes these layers for efficient feature extraction and pass them to classical layers, however, the quanvolution operation itself is not trainable. In order to enhance the adaptability of these layers within QuNNs it is important to make these layers trainable to harness their full potential. Making the quanvolutional layer trainable entails making the



**Fig. 5.** Detailed Methodology. A comprehensive overview of steps involved in the analysis of trainable QuNNs and proposed ResQuNNs. The results of quanvolutional layers are postprocess using both classical layer and quantum circuit. *Zero Padding* is performed to match the dimensions of input and output before passing it to residual block. The gradients accessibility through all the layers and training and validation accuracy are used as evaluation metrics. *QCL* quanvolutional layer.

quantum circuit (labeled as ❶ in Fig. 5) trainable. We do this by making the quanvolution layer as a pytorch layer which is later used as a standard layer during the model construction. The trainable QuNN utilized in this study is depicted in purple shaded region of Fig. 5, and operates through the following steps.

1. Encoding classical data into quantum states. We use a kernel of size  $2 \times 2$  which slides over the entire input image with stride 2. This results in 4 classical features extracted after every step, i.e.,  $x = \{x_1, x_2, x_3, x_4\}$ . These features are then encoded into quantum states, specifically into the rotation angles of *RY* gates. This approach is commonly called as angle encoding<sup>32</sup>. Each qubit  $|q_i\rangle$  is prepared in the ground state  $|0\rangle$ , and the *RY* gate applies a rotation by  $x_i$  radians on the Bloch sphere, which can be described by the following Equation:

$$U_{Enc_1} = |q_i\rangle = RY(x_i)|0\rangle \quad \text{where} \quad RY(\theta) = \begin{pmatrix} \cos(x/2) & -\sin(x/2) \\ \sin(x/2) & \cos(x/2) \end{pmatrix} \quad (1)$$

A total of 4 qubits will be required, for encoding 4 input features:

$$U_{Enc_1} = \bigotimes_{i=1}^4 RY(x_i)|0\rangle \quad (2)$$

2. Trainable quantum circuit operations. The encoded quantum states are then subjected to unitary evolution via an *RX* rotation gate. Let  $\theta_i$  be the angle of rotation for the  $i$ -th qubit applied by the *RX* gate:

$$U_{Rotation} = RX(\theta_i)|q_i\rangle \quad \text{where} \quad RX(\theta) = \begin{pmatrix} \cos(\theta/2) & -i \sin(\theta/2) \\ -i \sin(\theta/2) & \cos(\theta/2) \end{pmatrix}, \quad (3)$$

For 4-qubit layer, the above equation can written as:

$$U_{Rotation} = \bigotimes_{i=1}^4 RX(\theta_i) \quad (4)$$

To introduce entanglement between qubits in the quanvolutional layers, *CNOT* gates are applied between nearest neighbor qubits. Since we use total of 4 qubits, the entanglement operation in this context can be defined as:

$$U_{Entangle} = \bigotimes_{j=1}^3 CNOT_{j,j+1} \quad (5)$$

3. Overall transformation for quantum circuit. Combining the above steps, the overall transformation  $U_{Circuit_1}$  applied by the quanvolutional layer can be expressed as:

$$U_{\text{circuit}_1} = \left( \bigotimes_{i=1}^3 CNOT_{i,i+1} \right) \cdot \left( \bigotimes_{i=1}^4 RX(\theta_i) \right) \cdot \left( \bigotimes_{i=1}^4 RY(x_i) \right) \tag{6}$$

$$U_{\text{circuit}_1} = (U_{\text{Entangle}} \cdot U_{\text{Rotation}} \cdot U_{\text{Enc}_1}) \tag{7}$$

In the proposed architecture, the first quanvolutional layer (QCL1) is designed with a circuit depth of 4, while the second quanvolutional layer (QCL2) exhibits a reduced depth of 1. Both QCL1 and QCL2 incorporate identical quantum operations, consisting of single-qubit parameterized rotation gates and two-qubit CNOT entangling gates, as illustrated in label **●** of Fig. 5. The use of only single and two-qubits gates makes the quantum circuits compatible with NISQ devices. The distinction between the two layers, i.e., QCL1 and QCL2 lies solely in the number of repetitions of the quantum circuit, thereby differentiating their respective depths. This configuration (deep quantum layers in QCL1 and shallow layers in QCL2) is strategically chosen to distinctly evaluate the impact on performance when gradients from both layers are accessible in ResQuNNs, compared to the scenarios where the gradients of only the second layer are accessible. The reason is that when QCL1 is configured with a reduced depth, or when both QCL1 and QCL2 have equal depths, the effect of QCL1 on resulting performance difference may not be as pronounced or readily observable, even in cases where gradients from the second layer are accessible. Applied to the initial state  $|0000\rangle$ , the final state after the first QCL can be expressed as:

$$|\psi_{\text{QCL}_1}\rangle = \prod_{l=1}^4 U_{\text{circuit}_1} |0000\rangle \tag{8}$$

where  $l$  represents the total number of layers, i.e., number of times the circuit is repeated before measurement. The size of output after every quanvolutional layers is detailed in Table 2.  $|\psi_{\text{QCL}_1}\rangle$  is the final state state after the first quanvolutional layer before measurement and upon measurement the quantum state collapses and produce a classical result:

$$y = \mathcal{M}|\psi_{\text{QCL}_1}\rangle \tag{9}$$

where  $\mathcal{M}$  denotes the qubit measurement in computational basis. Since the result from first quanvolutional layer is classical, we need to encode again before passing it to the second quanvolutional layer using the same equation (Eq. 1):

$$U_{\text{Enc}_2} = |q_i\rangle = RY(y_i)|0\rangle \tag{10}$$

The kernel size remains the same so a total of 4 input features will be encoded from in every step:

$$y_{\text{enc}} = U_{\text{Enc}_2} = \bigotimes_{i=1}^4 RY(y_i) \tag{11}$$

The rest of the operations of trainable quantum circuit in QCL2 are same as above and the final circuit equation will be:

$$U_{\text{circuit}_2} = \left( \bigotimes_{i=1}^3 CNOT_{i,i+1} \right) \cdot \left( \bigotimes_{i=1}^4 RX(\theta_i) \right) \cdot \left( \bigotimes_{i=1}^4 RY(y_i) \right) \tag{12}$$

$$U_{\text{circuit}_2} = U_{\text{Entangle}} \cdot U_{\text{Rotation}} \cdot U_{\text{Enc}_2}$$

The depth of quantum circuit in QCL is set to 1. Hence, the final quantum state after the application of QCL2 will be:

Input size	Residual configuration	Output feature size
$28 \times 28 \times 1$	No residual	$7 \times 7 \times 1$
$28 \times 28 \times 1$	$X + O1$	$14 \times 14 \times 1$
$28 \times 28 \times 1$	$O1 + O2$	$14 \times 14 \times 1$
$28 \times 28 \times 1$	$X + O2$	$28 \times 28 \times 1$
$28 \times 28 \times 1$	$(X + O1) + O2$	$28 \times 28 \times 1$

**Table 2.** Output feature sizes for two quanvolutional layer.

$$|\psi_{\text{QCL}_2}\rangle = U_{\text{circuit}_2} |y_{\text{enc}}\rangle \quad (13)$$

$$\text{Output}_{\text{Quantum}} = \mathcal{M} |U_{\text{QCL}_2}\rangle \quad (14)$$

4. The measurement yields a list of classical expectation values. Similar to a classical convolution layer, each expectation value corresponds to a separate channel in a single output pixel.
5. The outputs from quantum layer are then postprocessed individually using both classical layer and quantum circuit.

#### Classical layer to postprocess quantum layer's output

First, we use classical layer to postprocess the output of quantum layer, as highlighted in label ② in Fig. 5). This is a major practice in state-of-the-art QuNN architectures, where a fully connected neuron layer is used to take the output of the quantum layer and produces the final classification output. The number of neurons in this layer are equal to the total number of classes in the dataset used, which in our case is 10 since we used MNIST handwritten digit dataset which has 10 classes (more details in “Experimental setup”).

Let's assume the output of the quantum layer is a set of measurements  $\{M = m_1, m_2 \dots m_n\}$ , where each  $m_i$  represents the measurement of a qubit in terms of classical information, such as the expectation values in computational basis. This result is transformed to a vector format and then passed to the fully connected layer with 10 neurons. The output feature size from quantum layers is different and depends on the kernel size and stride similar to that of classical convolutional layers. We experiment with a kernel size of  $2 \times 2$  and stride of 2. The output feature size after each quantum layer for different residual configuration (label ⑤ in Fig. 5) is presented in the Table 2.

The classical neural network layer with 10 neurons is a fully connected layer, where each neuron receives inputs from all  $n$  measurements of the quantum layer and processes these inputs to produce an output. The output of  $j$ -th neuron in the classical layer can be described by the following Equation:

$$y_j = f \left( \sum_{i=1}^n w_{ji} \cdot m_i + b_j \right)$$

where:

- $f$  is the SoftMax activation function, which converts a vector of  $K$  real numbers into a probability distribution of  $K$  possible outcomes. It is frequently used in multiclass classification problems.
- $w_{ji}$  is the weight from the  $i$ -th input to the  $j$ -th neuron.
- $b_j$  is the bias term for the  $j$ -th neuron.
- $y_j$  is the output of the  $j$ -th neuron.

The operation of this layer can be compactly represented in matrix form as:

$$Y = f(W \cdot M + B)$$

where:

- $Y$  is the output vector  $[y_1, y_2, \dots, y_{10}]$ .
- $f$  is applied element-wise to the result.
- $W$  is the weight matrix.
- $M$  is the vector of measurements from the quantum layer.
- $B$  is the bias vector.

#### Quantum circuit to postprocess quantum layer's output

We also use quantum circuit to postprocess the result from quantum layer, as highlighted by label ③ in Fig. 5. This is primarily for two reasons: firstly to demonstrate the working of fully quantum QuNNs, and secondly, to better understand the importance of residual connections by removing the influence of classical layers being used at the end. The use of a quantum circuit for post-processing introduces certain computational overheads, mainly due to the prerequisite of re-encoding the (classical) outputs from the quantum layers, a process that entails considerable computational expense. Also, the following trainable quantum circuit for postprocessing quantum layer's results after the encoding, takes longer to produce the results than a simple classical neuron layer. We use 4 classes from MNIST dataset for experimental simplification. For the purpose of encoding the outputs from the quantum layer, *Amplitude Encoding*<sup>33</sup>, is used which can be described by the following Equation:

$$|\psi_x\rangle = \sum_{i=1}^{2^n} x_i |i\rangle \quad (15)$$

where  $x_i$  is the  $i$ th element of  $x$  and  $|i\rangle$  is the  $i$ th computational basis state. We used amplitude encoding for postprocessing the results of quantum layers primarily because it allows to encode  $2^n$  features in  $n$  qubits<sup>32</sup>. Based on Table 2, the smallest output feature size from quantum layer(s) is  $7 \times 7 \times 1 = 49$ , which

requires at least 49 qubits via angle encoding (one qubit per feature) and only 6 qubits via amplitude encoding. Similarly, across the range of experiments we performed, the maximum output size of quanvolutional layer is  $28 \times 28 \times 1 = 784$ , requiring 784 qubits, if encoded in qubit rotation angles, whereas amplitude encoding reduces the requirement to merely 10 qubits. The use of 784 qubits is not practical in current NISQ devices, therefore we use amplitude encoding to encode the quanvolutional layer's result for further processing by a quantum circuit. To ensure a consistent impact of the quantum circuit's post-processing across all residual configurations, a 10-qubit circuit is utilized for processing the outputs of the quanvolutional layer, corresponding to the maximum feature size of 784. However, as mentioned earlier, we use 4 classes, therefore, only 4 qubits are measured each corresponding to one class. The output of quanvolutional, once encoded, is subjected to the unitary evolution through a simple quantum circuit, which can be described by the following Equation:

$$U(\theta) = \bigotimes_{j=1}^9 (CNOT_{j,j+1}) \cdot \bigotimes_{i=1}^{10} R_{y_i}(\theta) \quad (16)$$

### Proposed residual quanvolutional neural networks (ResQuNNs)

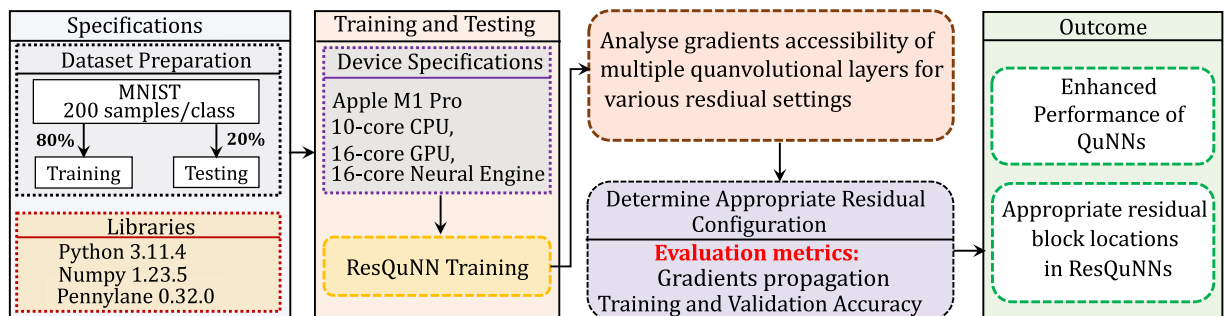
We observe that optimizing multiple (two or more) quanvolutional layers presents a significant challenges. One of the main difficulties lies in gradients accessibility across these layers during training, which are essential in gradient-driven optimization techniques. We typically notice that when we train two or more quanvolutional layers, the gradients of only the last layer are accessible. We have shown these results as a primary motivation of this paper in “Motivational case study”. To overcome the gradient accessibility issue encountered in multi-layered QuNNs, we introduce ResQuNNs. Drawing inspiration from classical residual neural networks, ResQuNNs incorporate residual blocks between quanvolutional layers. The residual blocks combine the output of existing quanvolutional layers with their own or the previous layer's input before forwarding it to the subsequent quanvolutional layer, as shown in label ④ in Fig. 5. These residual connections play a pivotal role in facilitating comprehensive gradient access and, in turn, improving the overall training performance of QuNNs. The integration of residual blocks ensures that gradients flow efficiently across the network, thereby overcoming the challenges associated with training multiple quanvolutional layers.

### Strategic placement of residual blocks

Within the proposed ResQuNN architecture, not every random residual configuration results in comprehensive gradient flow throughout the network, and the residual blocks needs to be strategically placed between quanvolutional layers. In our empirical investigation, we conduct extensive experimentation to determine the appropriate locations for inserting residual blocks within ResQuNNs comprising two quanvolutional layers, as shown in Fig. 5, labeled ⑤. This analysis helps us uncover the most effective residual configurations for ResQuNNs, enabling the gradients to flow throughout the proposed ResQuNN architecture, resulting in a deeper architecture and eventually the enhanced training performance. Based on the results of two quanvolutional layers, only the most effective configurations, i.e., those facilitating gradient flow through both layers were chosen for subsequent examination within three-layer quanvolutional frameworks, as denoted by label ⑥ in Fig. 5. We then trained ResQuNNs with two quanvolutional layers for all residual configurations for a multiclass classification problem (more on dataset in “Experimental setup”) to assess the impact of gradient accessibility, both throughout the entire network (all quanvolutional layers) and within segments of the network (only the last quanvolutional layer), on the learning efficacy of QuNNs.

### Experimental setup

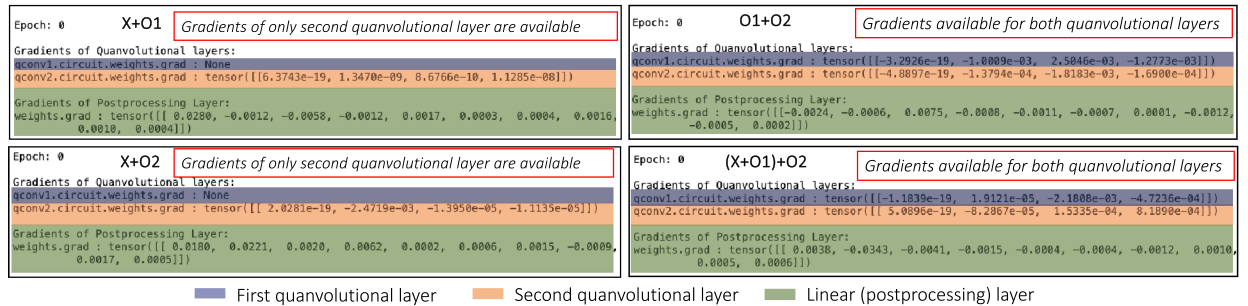
An overview of the experimental toolflow utilized in this paper is depicted in Fig. 6. We trained the proposed ResQuNNs on a selected subset of the MNIST dataset. Table 3 summarizes the dataset and hyperparameters specification used in this paper. The ResQuNNs are trained for 30 epochs. Adam optimizer with learning rate of 0.01 is used for the optimization. The objective function chosen for training is cross entropy loss, which can be described (for C classes) by the following equation:



**Fig. 6.** Overview of our experimental tool-flow.

Dataset specs				BS	Opt	LR	Quantvolutional layer specs			Postprocessing	
Total classes	Samples/class	Train data	Test data				Encoding	Kernel size	# of qubits	Classical	Quantum
10	200	1600	400	16	Adam	0.01	Angle (RY)	2 × 2	4	Dense layer (10 neurons)	–
4	200	640	160	16	Adam	0.01	Angle (RY)	2 × 2	4	–	10-Qubit circuit (4 measured)

**Table 3.** Dataset and hyperparameters specifications. *BS* batch size, *Opt* optimizer, *LR* learning rate.



**Fig. 7.** Gradients propagation through the network for two quantvolutional layers for different residual configurations. X = input, O1, O2 = output of first and second quantvolutional layer respectively.

$$H(p, q) = - \sum_{c=1}^C p(c) \log(q(c))$$

where  $H(p, q)$  is the cross entropy loss,  $p(c)$  represents the true probability distribution for class  $c$  and  $q(c)$  denotes the predicted probability distribution for class. All experiments are executed utilizing PennyLane, a Python library tailored for differentiable programming on quantum computing platforms<sup>34</sup>.

## Results and discussion

### Two quantvolutional layers

We use kernel size of  $2 \times 2$ , which results in 4-qubit quantvolutional layers. Our experimental framework involved extensive testing of all residual configurations as outlined in label ⑤ of Fig. 5. For each residual configuration, we investigated the gradient propagation through the quantvolutional layers. Afterwards, we trained the ResQuNNs for all these configurations to assess the impact of gradient accessibility throughout the network’s layers.

#### Analysis of gradients accessibility

In the case of *No Residual* setting (which refers to the configuration in which no residual blocks are incorporated within the multi-layered QuNNs and serves as a benchmark for comparison) we have already demonstrated in Fig. 3 that the gradients are exclusively accessible for the second quantvolutional layer. Therefore, we now proceed to investigate the gradient propagation through the quantvolutional layers for all other residual settings (from label ⑥ in Fig. 5), the results of which are illustrated in Fig. 7. It is worth mentioning here that while the gradient accessibility results are presented with quantum layer postprocessing, the accessibility of gradients remains the same across layers with classical postprocessing as well.

Our analysis reveals that specific residual configurations, notably  $O1 + O2$  and  $(X + O1) + O2$ , significantly enhance gradient propagation through the quantvolutional layers. This improvement not only makes them more conducive to gradient-based optimization strategies but also facilitates deep learning in QuNNs. In contrast, certain configurations, such as  $X + O1$  and  $X + O2$ , mirror the conditions of a *no residual* setting, where gradients fail to propagate past the terminal quantvolutional layer. Consequently, these configurations may exhibit a constrained training potential, impacting their effectiveness in QuNNs.

#### Analysis of training performance

We then trained the proposed ResQuNNs on a subset of MNIST dataset, as detailed in “Experimental setup”, for all the residual configurations labeled as ⑥ in Fig. 5, separately with classical and quantum postprocessing. The experimental procedures are consistent across all trials, adhering to the parameters discussed in “Experimental setup”. It is worth mentioning here that the *No residual* configuration also serves as a common benchmark for all the other residual configurations to analyze whether or not the inclusion of residual connections bring any advantage. Below, we separately discuss the training results of ResQuNNs both with classical and quantum postprocessing.


1. Classical postprocessing. We first evaluate the training performance of ResQuNNs that utilize classical neuron layers to postprocess the output of quantvolutional layers, as highlighted by label ② in Fig. 5. The training results of ResQuNNs with classical postprocessing are shown in Fig. 8. A common and a key observation is

that the incorporation of residual connections consistently improves the training performance of QuNNs, regardless of the specific residual configuration. For instance, while only the gradients of second quanvolutional layer are accessible in both *No residual* and  $X + O1$  configurations, the introduction of a residual connection in  $X + O1$  leads to superior performance compared to the *No residual* approach. We now individually explore the impact of different residual configurations on the performance of QuNNs with a focus on training efficiency and generalization capabilities (validation accuracy).

- **No residual.** When there is *No residual* block, the QuNNs exhibit *no training* as shown in label ❶ of Fig. 8. Although, even in case of no residual, the gradients of second quanvolutional layer are available but a significant information is lost in the first quanvolutional layer (firstly because of the  $2 \times 2$  quanvolution filter, and secondly due to the measurements and re-encodings), which does not contribute to the learning because of no gradients.
- **$X+O1$ .** When the residual block is added after the first quanvolutional layer, i.e., the input is added to the output of first quanvolutional and then passed as input to the second quanvolutional layer, the model shows significant performance enhancement with improved training and validation accuracy compared to *No residual* setting, as shown in label ❷ of Fig. 8. While  $X + O1$  residual configuration marks a significant improvement over the *No residual* configuration, it nonetheless presents a sub-optimal training performance as the first quanvolutional layer does not take part in the overall optimization because of no gradients.
- **$O1+O2$ .** Contrary to the *No residual* and  $X + O1$  configurations, the residual setup of  $O1 + O2$  facilitates the propagation of gradients through both quanvolutional layers, which is considered more conducive for gradient-based optimization techniques. This facilitation of gradient propagation results in a marked improvement in performance with the  $O1 + O2$  configuration, as depicted in label ❸ of Fig. 8. The training and generalization capabilities of the  $O1 + O2$  configuration significantly exceed those of  $X + O1$ , demonstrating approximately a 31% increase in training accuracy and a 45% enhancement in validation accuracy, respectively. This superiority can be attributed to the  $O1 + O2$  configuration's facilitation of gradient propagation through both quanvolutional layers, in contrast to  $X + O1$ , where gradient access is limited exclusively to the final quanvolutional layer.
- **$X+O2$  and  $(X+O1)+O2$ .** The residual configurations where the input is propagated directly ( $X + O2$ ) or indirectly ( $(X + O1) + O2$ ) to the output of second quanvolutional layer, the ResQuNNs outperform other residual configurations, as shown in and of Fig. 8. Additionally, both these residual settings exhibits very similar performance.

We argue that the residual configurations facilitating gradient flow through both quanvolutional layers, specifically,  $O1 + O2$  and  $(X + O1) + O2$ , will inherently exhibit superior performance in terms of both training and generalization compared to those configurations where gradient access is confined to the last quanvolutional layer only (*No Residual*,  $X + O1$ , and  $X + O2$  configurations). This is evident from the performance comparison of  $X + O1$  and  $O1 + O2$  (label ❷ and ❸ in Fig. 8), where  $O1 + O2$  performs significantly better. Additionally, we hypothesize that the superior performance of  $X + O2$  and  $(X + O1) + O2$  can be attributed to the last classical layer (used to meaningfully post process quanvolutional layer's output) in the network since the (input) classical data in both these settings is directly ( $X + O2$ ) or indirectly ( $(X + O1) + O2$ ) passed to the classical layer which tends to overshadow the learning contribution from quanvolutional layers. Therefore, in next section, we scrutinize the influence of classical layer (used for postprocessing) on the overall learning performance in all the residual configurations from Fig. 5 label ❹. This analysis is to identify the residual configurations yielding better training and generalization performance primarily due to the learning contributions from the quanvolutional layers.

**Comparison with benchmark models.** To assess the impact of the classical layer in various residual configurations of ResQuNNs, we conducted a detailed analysis using benchmark models. The benchmark models replicate the ResQuNNs with the residual configurations as shown in label ❹ of Fig. 5, with one key difference: the quanvolutional layers were made *untrainable*, i.e., their weights are not updated during training, and only the classical layers were subjected to training. This exercise allows us to determine the extent to which the classical layer tends to dominate the learning process, compared to the quanvolutional layers across all the residual configurations. To put it simply, configurations yielding similar performance when the quanvolutional layer remains untrained and only the classical layer is trained owe their success primarily to the classical layer. Conversely, configurations that demonstrate improved performance when both the quanvolutional and

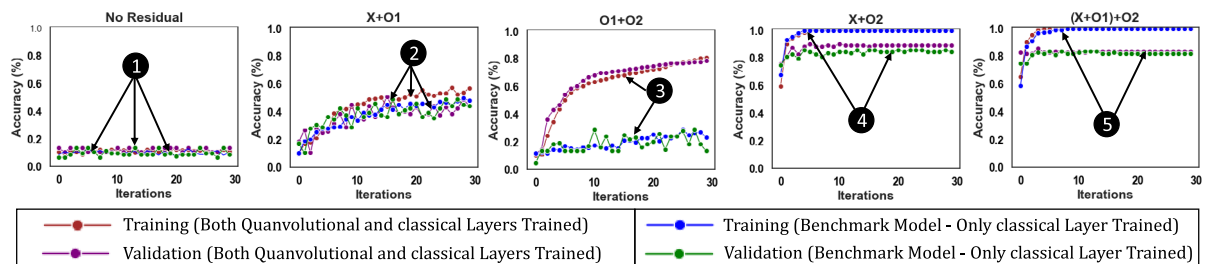


**Fig. 8.** Two quanvolutional layers—training results of all residual configurations with classical postprocessing and  $2 \times 2$  Kernel. No residual setting refers to multi-layered QuNNs without any residual block.


classical layers are trained, as opposed to when only the classical layers are trained, benefit from the residual configuration and enhanced gradient flow throughout the network.

The benchmark models are trained using the same hyperparameters and dataset as detailed in “Experimental setup”, and the results are compared with the models, where both quanvolutional and classical layers were trainable. The comparison results are presented in Fig. 9. The observed performance trends in these models gives insights into the effectiveness of quanvolutional layers in different residual configurations. *The key findings from this comparative analysis are summarized below:*


- **No Residual.** In the scenario where no residual connections are utilized, gradient flow is restricted exclusively to the second quanvolutional layer. Our observations indicate that the performance of the model, when both quanvolutional and classical layers undergo training, is very similar to that of the benchmark models in which training is restricted solely to the classical layers, with quanvolutional layers remaining static, as shown in label ❶ of Fig. 9. Nevertheless, in both settings, this particular configuration results in *minimal* or *No* training.
  - **Residual configuration  $X+O1$ .** In this residual configuration, despite gradient accessibility being limited to the second quanvolutional layer, the introduction of residual connections significantly enhances the model’s training capability. However, the comparative analysis of actual ResQuNNs (from Fig. 8) with benchmark models suggests that, the performance in this very residual configuration is predominantly influenced by the classical layer at the end of the quanvolutional layers, as highlighted by label ❷ in Fig. 9. This observation suggests that, within this configuration, the classical layer plays a critical role in the model’s overall learning and performance outcomes.
  - **Residual configuration  $(O1+O2)$ .** This residual configuration potentially enables the deep learning by permitting gradient propagation through the quanvolutional layers, thus ensuring their active involvement in the learning process. This results in a significant enhancement in the model’s performance, particularly when compared with the benchmark model. As denoted by label ❸ in Fig. 9, the performance of the model when both the quanvolutional layers along with classical layer are trained surpasses that of the benchmark model (where only classical layer is trained and quanvolutional layers are kept frozen) by approximately 66% in training accuracy and 64% in validation accuracy. This observation implies that in  $O1 + O2$  residual configuration, the learning performance is primarily driven by the quanvolutional layers. Such a notable enhancement underscores the benefit of fully utilizing both quanvolutional layers when gradients are accessible for them, highlighting the importance of their contribution to the learning process.
  - **Residual configuration with direct/indirect input to classical layer ( $X+O2$  and  $(X+O1)+O2$ ).** In scenarios where the input is either directly or indirectly routed to the classical layer at the network’s end, such as in  $X + O2$  and  $(X + O1) + O2$ , the performance of QuNNs with trainable quanvolutional demonstrates a close approximation to that of the benchmark models. This observation is highlighted by labels ❹ and ❺ in Fig. 9, indicating the predominant role of the classical layer in the learning process. The similar performance is achieved when only the classical layer is subjected to training, in comparison to scenarios where both quanvolutional and classical layers are trained, highlighting the significant role of the classical layer in the learning process in these residual configurations. However, considering the enhanced performance in the  $O1 + O2$  configuration, which allows gradient flow through both the quanvolutional layers, it suggests that  $(X + O1) + O2$  could be effective in scenarios without a classical layer at the end. This is because it also allows gradient propagation across all quanvolutional layers, potentially enabling the deep learning in QuNNs and eventually enhancing the learning performance, when the classical layer’s dominance is removed. To further elaborate on this, in the following section, we use the quantum circuit to postprocess the results of quanvolutional layers to further understand the role of residual approach in enabling deep learning in QuNNs, when there are no classical layers involved.
2. **Quantum postprocessing** In this section, we present the training results of ResQuNNs using quantum layer/circuit to postprocess the results of quanvolutional layer(s), as shown in label ❸ of Fig. 5. The primary reason behind quantum postprocessing is to reduce the influence of classical layer on the overall learning performance. While using quantum postprocessing, we observe clear advantage of residual configurations that allow the gradient flow across both the layers of the network compared to those which allow partial gradient



**Fig. 9.** Two quanvolutional layers—comparison of ResQuNNs with benchmark models for  $2 \times 2$  Kernel. No residual setting refers to multi-layered QuNNs without any residual block.



**Fig. 10.** Two quanvolutional layers—training results of all residual configurations with quantum postprocessing  $\times 2$  kernel. No residual setting refers to multi-layered QuNNs without any residual block.



**Fig. 11.** Gradients propagation through the network for three quanvolutional layers for (O1+O2)+O3 (left) and ((X+O1)+O2)+O3 (right).

flow. The training results are depicted in Fig. 10. Below, we separately discuss the results of ResQuNN with different residual configurations (from label 5 in Fig. 5), employing two quanvolutional layers.

- Configurations with restricted gradient flow (no residual, X+O1, X+O2). The residual configurations that allow partial gradient flow, i.e. *No residual*, X+O1 and X+O2 exhibits no to negligible training, as shown in labels 1, 2 and 4 of Fig. 10. This outcome underscores the critical role of gradient flow in the training process of QuNNs.
- Configurations facilitating full gradient flow (X+O2, (X+O1)+O2). The residual configurations that facilitate gradients flow through all the layers of the network specifically ‘X + O2’ and ‘(X + O1) + O2’, demonstrated significantly enhanced performance and completely outperforms other residual configurations, as illustrated by label 3 and 5 in Fig. 10.

The contrast in training performance across different residual configurations elucidates the tremendous importance of gradient flow in the architecture of ResQuNNs. Configurations enabling full gradient availability across network layers emerged as distinctly advantageous, significantly improving the training potential of ResQuNNs. Conversely, configurations with limited gradient flow manifested negligible training outcomes, highlighting the inadequacy of partial gradient availability for effective network learning. This distinction underlines the significance of architectural choices in the design of ResQuNNs, particularly in the context of integrating quantum and classical layers for enhanced learning capabilities.

### Three quanvolutional layers

In previous sections, we established the critical importance of gradient accessibility across all the layers in QuNNs for achieving better training performance. In this section, we extend our exploration to demonstrate that our proposed ResQuNNs can further facilitate the utilization of an increased number of quanvolutional layers for more deeper networks. Specifically, we consider the scenario of training QuNNs with three quanvolutional layers. Due to space constraints, we only present the results for residual configurations where gradients are accessible across all three quanvolutional layers, as shown in Fig. 11. Out of 15 potential residual configurations examined, only two, specifically (O1 + O2) + O3 and ((X + O1) + O2) + O3, allowed gradient access across all three layers. The other configurations restricted gradient flow to one or two layers. This further underscores the significance of our proposed residual approach, as it enables the desired comprehensive gradient accessibility, with greater number of layers, effectively enabling the deep learning in QuNNs, which can be pivotal for optimizing complex problems.

### Conclusion

In this paper, we presented a framework for developing Quanvolutional Neural Networks (QuNNs) by introducing trainable quanvolutional layers and the concept of Residual Quanvolutional Neural Networks (ResQuNNs) for enabling the deep learning in QuNNs. Our approach addresses the significant challenges of gradient flow and adaptability in deep QuNN architectures, enabling enhanced training performance and scalability. A key contribution of this work is the development of trainable quanvolutional layers, which enhance the learning potential of QuNNs, and the strategic incorporation of residual blocks to ensure comprehensive gradient accessibility across the network in multi-layered QuNNs, thus, enabling deeper, more complex QuNNs architectures to be effectively trained. Our empirical findings highlight the importance of optimal placement of residual blocks, providing critical insights for future QuNN architecture designs with even more deeper

layers. In summary, our work signifies a substantial leap in quantum deep learning, offering new directions for theoretical and practical quantum computing research. By tackling core issues of adaptability and gradient flow, our work paves the way for the development of more robust and efficient quantum neural networks, marking a pivotal step toward harnessing the full potential of quantum technologies.

### Data availability

The datasets used and/or analysed during the current study available from the corresponding author on reasonable request.

Received: 8 September 2024; Accepted: 4 June 2025

Published online: 01 July 2025

### References

- Schuld, M., Sinayskiy, I. & Petruccione, F. An introduction to quantum machine learning. *Contemp. Phys.* **56**, 172–185. <https://doi.org/10.1080/00107514.2014.964942> (2014).
- Biamonte, J. et al. Quantum machine learning. *Nature* **549**, 195–202. <https://doi.org/10.1038/nature23474> (2017).
- Kashif, M., Marchisio, A. & Shafique, M. Computational advantage in hybrid quantum neural networks: Myth or reality? [arXiv:2412.04991](https://arxiv.org/abs/2412.04991) (2025).
- Kashif, M. & Al-Kuwari, S. Demonstrating quantum advantage in hybrid quantum neural networks for model capacity. In *2022 IEEE International Conference on Rebooting Computing (ICRC)*. 36–44. <https://doi.org/10.1109/ICRC57508.2022.00011> (2022).
- Wei, S. et al. A quantum convolutional neural network on NISQ devices. *AAPPS Bull.* **32**. <https://doi.org/10.1007/s43673-021-00030-3> (2022).
- Zaman, K. et al. A survey on quantum machine learning: Current trends, challenges, opportunities, and the road ahead (2023). [arXiv:2310.10315](https://arxiv.org/abs/2310.10315).
- Cong, I. et al. Quantum convolutional neural networks. *Nat. Phys.* **15**, 1273–1278. <https://doi.org/10.1038/s41567-019-0648-8> (2019).
- Rajesh, V. et al. Quantum convolutional neural networks (QCNN) using deep learning for computer vision applications. In *2021 RTEICT*. 728–734. <https://doi.org/10.1109/RTEICT52294.2021.9574030> (2021).
- Hur, T. et al. Quantum convolutional neural network for classical data classification. *Quantum Mach. Intell.* **4**. <https://doi.org/10.1007/s42484-021-00061-x> (2022).
- Cerezo, M., Verdon, G., Huang, H.-Y., Cincio, L. & Coles, P. J. Challenges and opportunities in quantum machine learning. *Nat. Comput. Sci.* **2**, 567–576. <https://doi.org/10.1038/s43588-022-00311-3> (2022).
- Shen, G. et al. Quantum convolutional neural network for image classification. *Pattern Anal. Appl.* **26**. <https://doi.org/10.1007/s10044-022-01113-z> (2023).
- Kashif, M., Rashid, M., Al-Kuwari, S. & Shafique, M. Alleviating barren plateaus in parameterized quantum machine learning circuits: Investigating advanced parameter initialization strategies. In *2024 Design, Automation & Test in Europe Conference & Exhibition (DATE)*. 1–6. <https://doi.org/10.23919/DATED58400.2024.10546644> (2024).
- Schuld, M. & Killoran, N. Quantum machine learning in feature Hilbert spaces. *Phys. Rev. Lett.* **122**. <https://doi.org/10.1103/physrevlett.122.040504> (2019).
- Sebastian, A. et al. Beyond qubits : An extensive noise analysis for qutrit quantum teleportation (2023). [arXiv:2309.12163](https://arxiv.org/abs/2309.12163).
- Hu, M. et al. Strong quantum nonlocality with genuine entanglement in an  $n$ -qutrit system (2023). [arXiv:2308.16409](https://arxiv.org/abs/2308.16409).
- Kim, J. et al. Classical-to-quantum convolutional neural network transfer learning. *Neurocomputing* **555**. <https://doi.org/10.1016/j.neucom.2023.126643> (2023).
- Henderson, M. et al. Quconvolutional neural networks: Powering image recognition with quantum circuits. *Quantum Mach. Intell.* **2**. <https://doi.org/10.1007/s42484-020-00012-y> (2020).
- Matic, A. et al. Quantum-classical convolutional neural networks in radiological image classification. In *2022 IEEE International Conference on Quantum Computing and Engineering (QCE)*. 56–66. <https://doi.org/10.1109/QCE53715.2022.00024> (IEEE Computer Society, 2022).
- Chen, S. et al. Quantum convolutional neural networks for high energy physics data analysis. *Phys. Rev. Res.* **4**. <https://doi.org/10.1103/PhysRevResearch.4.013231> (2022).
- Meedinti, G. et al. A quantum convolutional neural network approach for object detection and classification (2023). [arXiv:2307.08204](https://arxiv.org/abs/2307.08204).
- Baek, H. et al. Scalable quantum convolutional neural networks (2022). [arXiv:2209.12372](https://arxiv.org/abs/2209.12372).
- Matondo-Mvula, N. & Elleithy, K. Breast cancer detection with quantum convolutional neural networks. *Entropy* **26**, 630 (2024).
- Ahalya, R., Almutairi, F. M., Snehalatha, U., Dhanraj, V. & Aslam, S. M. Ranet: A custom cnn model and quantum convolutional neural network for the automated detection of rheumatoid arthritis in hand thermal images. *Sci. Rep.* **13**, 15638 (2023).
- Sebastianelli, A., Mauro, F. et al. Quanv4eo: Empowering earth observation by means of quantum convolutional neural networks. [arXiv:2407.17108](https://arxiv.org/abs/2407.17108) (2024).
- Maouaki, W. E., Marchisio, A., Said, T., Shafique, M. & Bennai, M. Robqunns: A methodology for robust quantum convolutional neural networks against adversarial attacks. [arXiv:2407.03875](https://arxiv.org/abs/2407.03875) (2024).
- Maouaki, W. E., Marchisio, A., Said, T., Bennai, M. & Shafique, M. Advqunns: A methodology for analyzing the adversarial robustness of quantum convolutional neural networks. [arXiv preprint arXiv:2403.05596](https://arxiv.org/abs/2403.05596) (2024).
- Vu, T. H., Le, L. H. & Pham, T. B. Exploring the features of quantum convolutional neural networks for improved image classification. *Quantum Mach. Intell.* **6**, 29 (2024).
- Bhatia, A. S., Kais, S. & Alam, M. A. Federated quantum convolutional neural network: A new paradigm for collaborative quantum learning. *Quantum Sci. Technol.* **8**, 045032 (2023).
- Mattern, D., Martyniuk, D., Willems, H., Bergmann, F. & Paschke, A. Variational quantum convolutional neural networks with enhanced image encoding. [arXiv:2106.07327](https://arxiv.org/abs/2106.07327) (2021).
- Kashif, M. & Al-Kuwari, S. Resqunets: A residual approach for mitigating barren plateaus in quantum neural networks. *EPJ Quant. Tech.* **11**. <https://doi.org/10.1140/epjqt/s40507-023-00216-8> (2024).
- Wen, J., Huang, Z., Cai, D. & Qian, L. Enhancing the expressivity of quantum neural networks with residual connections. *Commun. Phys.* **7**, 220 (2024).
- Kashif, M. & Al-Kuwari, S. Design space exploration of hybrid quantum-classical neural networks. *Electronics* **10**, 2980 (2021).
- LaRose, R. & Coyle, B. Robust data encodings for quantum classifiers. *Phys. Rev. A* **102**. <https://doi.org/10.1103/physreva.102.032420> (2020).
- Bergtholm, V. et al. PennyLane: Automatic differentiation of hybrid quantum-classical computations. [arXiv https://doi.org/10.48550/ARXIV.1811.04968](https://arxiv.org/abs/1811.04968) (2018).

## Acknowledgements

This work was supported in part by the NYUAD Center for Quantum and Topological Systems (CQTS), funded by Tamkeen under the NYUAD Research Institute grant CG008.

## Author contributions

M.K. and M.S. conceptualized and refined the research idea. M.K. implemented the idea, conducted the experiments, and generated the results, with M.S. providing supervision throughout the experimentation and implementation process. M.K. prepared the initial draft of the manuscript, while M.S. critically reviewed and edited the final version.

## Declarations

### Competing interests

The authors declare no competing interests.

## Additional information

**Correspondence** and requests for materials should be addressed to M.K.

**Reprints and permissions information** is available at [www.nature.com/reprints](http://www.nature.com/reprints).

**Publisher's note** Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.

**Open Access** This article is licensed under a Creative Commons Attribution-NonCommercial-NoDerivatives 4.0 International License, which permits any non-commercial use, sharing, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if you modified the licensed material. You do not have permission under this licence to share adapted material derived from this article or parts of it. The images or other third party material in this article are included in the article's Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit <http://creativecommons.org/licenses/by-nc-nd/4.0/>.

© The Author(s) 2025