

Differentiable Quantum Residual Networks without Post-Selection

Amena Khatun¹, Akib Karim¹ and Muhammad Usman^{1,2}

¹ Quantum Systems, Data61, CSIRO, Australia

² School of Physics, The University of Melbourne, Victoria, Australia
{amena.khatun, akib.karim, muhammad.usman}@csiro.au

Residual learning transformed the training of deep classical neural networks, yet its adoption in quantum machine learning has remained limited because existing quantum residual architectures rely on post-selection, which prevents fully differentiable optimisation and makes execution increasingly inefficient as circuit depth grows. We introduce a new quantum machine learning architecture, the post-selection-free Quantum Residual Neural Network (QResNet), which implements differentiable residual learning through deterministic ancilla-controlled residual channels. Unlike previous QResNet formulations, the proposed architecture eliminates the sample-and-discard step while retaining the optimisation benefits of residual learning, enabling end-to-end gradient-based optimisation of both variational circuit parameters and residual strengths. We further characterise the trainability of the architecture and show that its optimisation behaviour is governed jointly by the residual depth and the depth of the variational blocks, identifying a polylogarithmic regime that avoids barren plateaus while remaining beyond the regime of provably classically simulable shallow circuits. On the ten-class MNIST benchmark, the proposed architecture achieves comparable classification accuracy using approximately six times fewer quantum gates than a 200-layer quantum variational classifier (QVC); adding only five residual blocks to a 30-layer variational backbone improves accuracy from 65% to 80%, while also demonstrating improved resistance to black-box transfer attacks. These results establish differentiable quantum residual learning as a practical architectural primitive for trainable, hardware-efficient, and robust quantum machine learning.

1 Introduction

Quantum variational classifiers (QVCs) are the dominant framework for quantum machine learning (QML) on noisy intermediate-scale quantum (NISQ) devices. Despite significant advancement of QVCs, their practical deployment remains limited by barren plateaus, circuit depths incompatible with near-term hardware, and classical simulability [1–3]. Developing trainable and hardware-efficient QML architectures therefore remains an important open challenge. Residual connections have become a standard component of modern deep learning, enabling stable optimisation and gradient propagation in very deep neural networks. Recent efforts have introduced residual learning into QML through post-selected linear-combination-of-unitaries (LCU) frameworks [4], in which an ancilla prepared in superposition is post-selected to realise a weighted combination of the identity and a variational unitary. However, post-selection introduces probabilistic execution with acceptance probabilities that decay exponentially with circuit depth, and breaks end-to-end differentiability through its sample-and-discard procedure. Density-based quantum learning frameworks eliminate post-selection [5], but replace coherent residual learning with probabilistic mixtures and do not provide trainable control over information flow across successive layers.

In this work, we introduce a new quantum machine learning architecture that enables differentiable residual learning without post-selection. The proposed QResNet replaces post-selected residual operations with deterministic ancilla-controlled residual channels, eliminating the sample-and-discard step while enabling end-to-end gradient-based optimisation. Each ancilla is prepared, used as a control, and uncomputed without measurement, yielding a completely positive trace-preserving residual channel together with a deterministic surrogate objective for optimisation. The architecture jointly optimises the variational circuit parameters and trainable residual strengths. We further characterise the trainability of the proposed architecture, identify conditions under which barren plateaus can be avoided, and validate its performance on MNIST, CIFAR-2, and the real-world SARFish remote-sensing dataset.

2 Post-Selection-Free Quantum Residual Neural Network

The proposed QResNet implements residual learning through a sequence of ancilla-controlled residual blocks, illustrated in Fig. 1. Each residual block contains an ancilla qubit, a variational unitary $W_l(\vartheta_l)$, and a trainable residual-strength parameter β_l . The ancilla controls whether the variational transformation acts on the data register or whether the input state bypasses the transformation through an identity path, analogous to skip connections in classical residual networks. The key distinction from previous QResNet formulations is that the proposed architecture replaces post-selected coherent residual maps with deterministic, differentiable residual channels, allowing standard gradient-based optimisation without stochastic acceptance steps.




Fig. 1. Overview of the proposed post-selection-free QResNet architecture. (a) Sequence of ancilla-controlled residual blocks. (b) Variational residual block. (c) Binary-classification architecture. (d) Multi-class architecture combining a variational backbone and QResNet.

This prepare, control and uncompute sequence is related to the LCU technique. In LCU, an ancilla prepared in a superposition coherently selects between different unitaries, and post-selecting the ancilla outcome implements a linear combination of those unitaries on the data qubit. In Ref. [4], the ancilla amplitudes are chosen such that post-selection yields a residual map proportional to $I + \beta_l W_l$. In our approach, we adopt a different ancilla preparation rule. Tracing out the ancilla after the prepare-control-uncompute sequence yields a completely positive, trace-preserving channel on the data register. Let ρ denote the density matrix of the input state to the ℓ -th residual block,

$$\mathcal{E}_l(\rho) = \frac{1}{1 + |\beta_l|^2} \rho + \frac{|\beta_l|^2}{1 + |\beta_l|^2} W_l(\vartheta_l) \rho W_l(\vartheta_l)^\dagger, \quad (1)$$

a probabilistic mixture of the identity channel and the unitary conjugation by $W_l(\vartheta_l)$. The residual connections are thus modified from the original proposal: successive layers are combined by a classical mixture rather than a coherent superposition. Unlike the original QResNet formulation [4], which employs post-selection to realize a coherent residual map, we adopt an alternative ancilla preparation and trace out the ancilla without measurement. This yields the deterministic channel in Eq. (1) together with the surrogate objective in Eq. (2), enabling end-to-end differentiable optimisation.

$$f(x) = \left[\prod_{\ell=1}^L (1 + \beta_\ell^2) \right] \langle Z_0 \rangle. \quad (2)$$

Table 1. Performance comparison between standard QVC architecture and the proposed QResNet models across multiple datasets. Total gate counts are computed for 10 data qubits. For QVC-200, the classification accuracy is reported from Ref. [6]. The gates counts are calculated from the circuit structure.

Model	Qubits	Layers	Dataset	Task	Test Acc. (%)	Total Gates
QResNet (5 blocks)	10	5	MNIST	Binary	99	200
QResNet (5 blocks)	10	5	CIFAR-2	Binary	76	200
QResNet (5 blocks)	10	5	SARFish	Binary	72.14	200
QVC-200	10	200	MNIST	Multi-class (10)	85	8000
QVC-30	10	30	MNIST	Multi-class (10)	65	1200
QVC-30 + QResNet	10	30 + 5	MNIST	Multi-class (10)	80	1400

The resulting model is fully differentiable with respect to both the variational parameters and the residual strengths. The parameter β_l has a clear interpretation as the relative contribution of the identity and variational branches within the residual block. In the limit $\beta_l \rightarrow 0$, the block reduces to the identity operation and is effectively bypassed. Conversely, as $|\beta_l| \rightarrow 1$, the channel approaches an equal mixture of the identity transformation and the variational unitary channel. This parameterisation eliminates post-selection while retaining the optimisation benefits of residual learning, enabling the network to adaptively regulate information flow across successive layers through gradient-based optimisation.

3 Trainability and Experimental Results

The key observation is that the gradient of QResNet can be expressed as a weighted sum of multiple shallow variational pathways created by the residual architecture. Expanding the gradient as the number of qubits and residual blocks grows, each term in this sum corresponds to a different combination of residual blocks, so the trainability of the model depends on two quantities: the depth of the variational layers within each block and the total number of blocks. Building on known depth–trainability results for variational circuits [1], our analysis shows that barren plateaus are avoided when both the per-block depth and the number of blocks scale polylogarithmically with the qubit count. In this regime, QResNet is more expressive than a single QVC of comparable depth, because the residual construction effectively combines many shallow variational pathways. Comparable accuracy can therefore be reached with far fewer layers than a standard QVC, improving trainability while keeping circuits shallow. This polylogarithmic regime also lies outside the class of shallow circuits known to be classically simulable.

Table 1 summarises the results. For 10-class MNIST classification, a QVC–QResNet architecture achieves 80% accuracy using 1400 quantum gates, compared with approximately 85% accuracy from a 200-layer QVC requiring around 8000 gates. Adding only five residual blocks improves the performance of a 30-layer QVC baseline from 65% to 80% while increasing the gate count from 1200 to 1400, directly demonstrating the contribution of residual learning. The same five-block QResNet circuit (200 gates) was further evaluated on binary CIFAR-2 and the real-world SARFish remote-sensing dataset, achieving 76% and 72.14% accuracy, respectively, and demonstrating applicability beyond MNIST to higher-variability and noisy real-world data.

We additionally evaluate robustness against adversarial perturbations using FGSM attacks. While accuracy degrades under white-box attacks, the architecture shows improved resistance to black-box transfer attacks generated from a classical neural network. The reduced transferability indicates that QResNet learns decision boundaries that differ substantially from those of classical models, consistent with improved robustness under transferred attacks.

4 Conclusion

We presented a new differentiable quantum residual architecture that eliminates post-selection by implementing residual learning through deterministic ancilla-controlled residual channels. The proposed QResNet combines theoretical trainability guarantees with competitive experimental performance, identifying a polylogarithmic depth regime that avoids barren plateaus while achieving comparable accuracy using substantially fewer quantum gates than deep variational quantum classifiers. These results establish differentiable quantum residual learning as a practical architectural foundation for scalable, trainable, and hardware-efficient quantum machine learning.

Preprint: A full version of this work is available at arXiv:2604.06866

References

1. Martin Larocca, Supanut Thanasilp, Samson Wang, Kunal Sharma, Jacob Biamonte, Patrick J Coles, Lukasz Cincio, Jarrod R McClean, Zoë Holmes, and Marco Cerezo. Barren plateaus in variational quantum computing. Nature Reviews Physics, pages 1–16, 2025.
2. Supanut Thanasilp, Samson Wang, Nhat Anh Nghiem, Patrick Coles, and Marco Cerezo. Subtleties in the trainability of quantum machine learning models. Quantum Machine Intelligence, 5(1):21, 2023.
3. Valentin Heyraud, Zejian Li, Kaelan Donatella, Alexandre Le Boité, and Cristiano Ciuti. Efficient estimation of trainability for variational quantum circuits. PRX Quantum, 4(4):040335, 2023.
4. Jamie Heredge, Maxwell West, Lloyd Hollenberg, and Martin Sevier. Nonunitary quantum machine learning. Phys. Rev. Appl., 23:044046, Apr 2025.
5. Brian Coyle, Snehal Raj, Natansh Mathur, El Amine Cherrat, Nishant Jain, Skander Kazdaghli, and Iordanis Kerendis. Training-efficient density quantum machine learning. npj Quantum Information, 11(1):172, Nov 2025.
6. Maxwell T West, Sarah M Erfani, Christopher Leckie, Martin Sevier, Lloyd CL Hollenberg, and Muhammad Usman. Benchmarking adversarially robust quantum machine learning at scale. Physical Review Research, 5(2):023186, 2023.