Information, Computing and Intelligent Systems, 2025, No. 7, 49 – 60
UDC 004.415.2 https://doi.org/10.20535/2786-8729.7.2025.333887

COMPARATIVE ANALYSIS OF LCNet050 AND MobileNetV3 ARCHITECTURES IN HYBRID QUANTUM–CLASSICAL NEURAL NETWORKS FOR IMAGE CLASSIFICATION

Arsenii Khmelnytskyi*
https://orcid.org/0009-0009-1631-8172

Yuri Gordienko
https://orcid.org/0000-0003-2682-4668

National Technical University of Ukraine
"Igor Sikorsky Kyiv Polytechnic Institute", Kyiv, Ukrainee

*Corresponding author: ars4899@gmail.com

This study explores the impact of classical backbone architecture on the performance of hybrid quantum-classical neural networks in image classification tasks. Hybrid models combine the representational power of classical deep learning with the potential advantages of quantum computation. Specifically, this research employs a quanvolutional neural network architecture in which a quantum convolutional layer, based on a four-qubit Ry circuit, preprocesses input images before classical processing.

Despite the growing interest in hybrid models, few studies have systematically investigated how variations in classical architecture design affect the overall performance of hybrid quantum-classical neural networks. To address this gap, we compare two lightweight convolutional backbones – `MobileNetV3Small050` and `LCNet050` – integrated with an identical quantum preprocessing layer. Both models are evaluated on the `CIFAR-10` dataset using 5-fold stratified cross-validation. Performance is assessed using multiple metrics, including accuracy, macro- and micro-averaged area under the curve, and class-wise confusion matrices.

The results indicate that the `LCNet`-based hybrid model consistently outperforms its MobileNet counterpart, achieving higher overall accuracy and area under the curve scores, along with improved class balance and robustness in distinguishing less-represented classes. These findings underscore the critical role of classical backbone selection in hybrid quantum-classical architectures. While the quantum layer remains fixed, the synergy between quantum preprocessing and classical feature extraction significantly affects model performance. This study contributes to a growing body of work on quantum-enhanced learning systems by demonstrating the importance of classical design choices. Future research may extend these insights to alternative datasets, deeper or transformer-based backbones, and more expressive quantum circuits.

Keywords: neural networks, quantum computing, hybrid neural networks, image classification.

1. Introduction

Machine learning, and in particular Deep Learning (DL), remain active and rapidly evolving fields that have transformed numerous areas of science, technology, and society, such as computer vision.

Classical DL models have achieved high accuracy for various image classification tasks because they can learn and extract complex patterns from high-dimensional data [1]. However, as the size of data sets and computational requirements increase, classical models also have drawbacks, mainly in terms of scalability and power consumption.

These limitations have fueled interest in alternative models of computing, with quantum computing as a potential prospect. Quantum computing can exploit quantum mechanical effects, such as superposition and entanglement, to solve problems that cannot be solved by any conceivable classical computer [2]. In this sense, the Hybrid (quantum-classical) Neural Networks (HNNs) have attracted considerable attention as an empirical way to obtain the best possible performance from the current generation of quantum machines.

HNNs exploit the combined expressive power of classical HNNs with the computational capabilities of quantum algorithms [3]. Recent advances, such as quantum variational schemes and quantum convolution operations, have shown promise in image feature extraction and classification for complex visual recognition. However, the field is still in its infancy with a number of outstanding research challenges, including hardware scalability, quantum noise, and a poor theoretical understanding of the behavior of quantum models.

Given the rapid advancement of quantum computing and its increasing integration with classical machine learning systems, research in HNNs is of growing importance. Understanding how classical architectures affect performance in such systems is essential for unlocking their practical potential and guiding future developments in quantum-enhanced AI applications.

2. Literature review and problem statement

The past few years have seen growing research emphasis in the field of quantum computing and machine learning, specifically to construct hybrid quantum-classical models for image classification and pattern recognition. There have been a number of works on Quantum Variational Classifiers [4], quantum kernel methods [5], and quantum convolutional networks [6] that try to harness the unique features of quantum systems for enhancing learning power and model expressiveness.

Early on, hybrid models incorporated small quantum circuits into classical networks, typically by replacing or supplementing fully connected layers with quantum variational circuits. These works showed the potential of hybrid models on small datasets `MNIST`, `Fashion-MNIST`, and `CIFAR-10` in simulated or under-constrained qubit hardware. Particularly, works such as those of showed that these hybrid quantum–classical models were indeed capable of outperforming their solely classical counterparts in regions of parameter space where classical models failed to generalize.

Another very interesting field of research in this context has been the ideas of quantum convolutional layers – also known as quanvolutional layers – that do their best to approximate the behavior of classical convolutions while using quantum operations. These layers have empirically improved the ability to learn the features, especially when considering complex image datasets, and have been leveraged to form HNNs of different structures [7].

Nevertheless, a lack still exists in systematically investigating the interaction among quantum layers and different classical-based backbone architectures. However, from the existing works, it remains unknown how the architectural properties (e.g., depth, channel width, layer connectivity, etc.) affect the hybrid models. Among them, lightweight and mobile-compatibility networks, such as `MobileNetV3` and `LCNet` [8], are extensively used in classical settings but largely unexploited in quantum-classical hybrid circuits.

In summary, while prior studies demonstrate the potential of hybrid quantum-classical neural networks, they rarely explore the influence of specific classical backbone architectures within a controlled quantum framework. This gap in the literature–namely, the lack of systematic comparison between lightweight convolutional models such as `LCNet050` and `MobileNetV3Small050` in hybrid settings-suggests that further investigation is warranted. Therefore, it is reasonable to conduct a study dedicated to evaluating the impact of classical backbone selection on the performance and stability of hybrid quantum-classical models in image classification tasks.

3. The aim and objectives of the study

The aim of this study is to evaluate the effect of classical lightweight backbone architecture on the performance of HNNs for image classification. This is grounded in the identified research gap concerning the role of classical model selection in the overall effectiveness and stability of hybrid systems.

To achieve this aim, the following objectives are defined:

– to implement and train two hybrid neural network architectures that integrate a fixed quantum preprocessing layer with either the `LCNet050` or `MobileNetV3Small050` backbone;

– to assess and compare their classification performance using metrics such as *macro-averaged AUC* and *micro-averaged AUC*, *class-wise AUC*, and *confusion matrices* on the `CIFAR-10` dataset;

– to analyze the impact of backbone selection on model stability and generalization across multiple validation folds.

4. Materials and methods of investigating hybrid quantum-classical neural networks

The object of the study is the classification performance of HNNs with different classical backbones. To evaluate the impact of classical backbone architecture on the performance of HNN [9, 10], two model configurations were developed. Both integrate a quantum processing layer with a lightweight convolutional neural network, differing only in the classical backbone: one utilizes `MobileNetV3Small050`, while the other uses `LCNet050`.

This section presents the architecture of the hybrids and describes the quantum and classical components independently for clarity.

4.1. Evaluation metrics

A few performance measures were defined to evaluate the performance of the proposed hybrid structures:

– validation accuracy;
– validation *AUC (macro and micro)*;
– validation loss;
– classical model complexity.

It should be noted that for the quantum layers, we performed simulations on classical hardware and execution time was not utilized as a measure of performance. Instead, we recorded the number of quantum circuit evaluations in the training and testing phases to provide an estimate of the overhead involved in deploying the models to an actual quantum computer.

4.2. Quantum circuits

In this work, the proposed HNNs employ a 4-qubit *Ry-based quantum* circuit. Each qubit in the circuit comes with one trainable parameter and is initially prepared with a *Hadamard gate*, then with a *Ry rotation gate* [11].

Note that, in this case, the number of circuit inputs and outputs is the same and matches the number of qubits used as being illustrated in Fig. 1. This particular quantum circuit design was chosen following the results from our previous study, where it proved best in its flexibility and performance for hybrid neural network models.

4.3. Datasets

To demonstrate the efficiency of the proposed hybrid quantum-classical neural network architectures, the `CIFAR-10` dataset was considered. `CIFAR-10` is a commonly used benchmark in the computer vision field, especially for image classification. The dataset is divided into 10 object class categories: airplane, automobile, bird, cat, deer, dog, frog, horse, ship, and truck; containing a total of 60,000 color images of 32×32 pixels. Fig. 2 shows a section of the `CIFAR-10` dataset.

Out of the entire dataset, 50,000 images is used for training and the remaining 10,000 images for testing. A validation set from the training data was used to keep track of performance and overfitting during the training. All images were normalized to zero mean and unit variance for each channel, and standard augmentations including random horizontal flipping and cropping were applied to improve the generalizability of the models. These preprocessing steps ensured consistency across training runs and helped reduce noise-related variance.

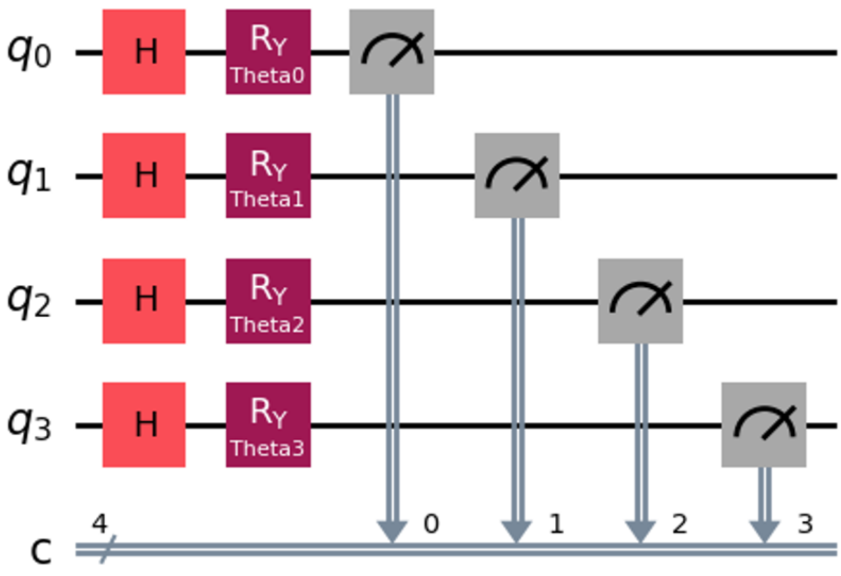

Fig. 1. Diagram of *Ry quantum* circuit used in research

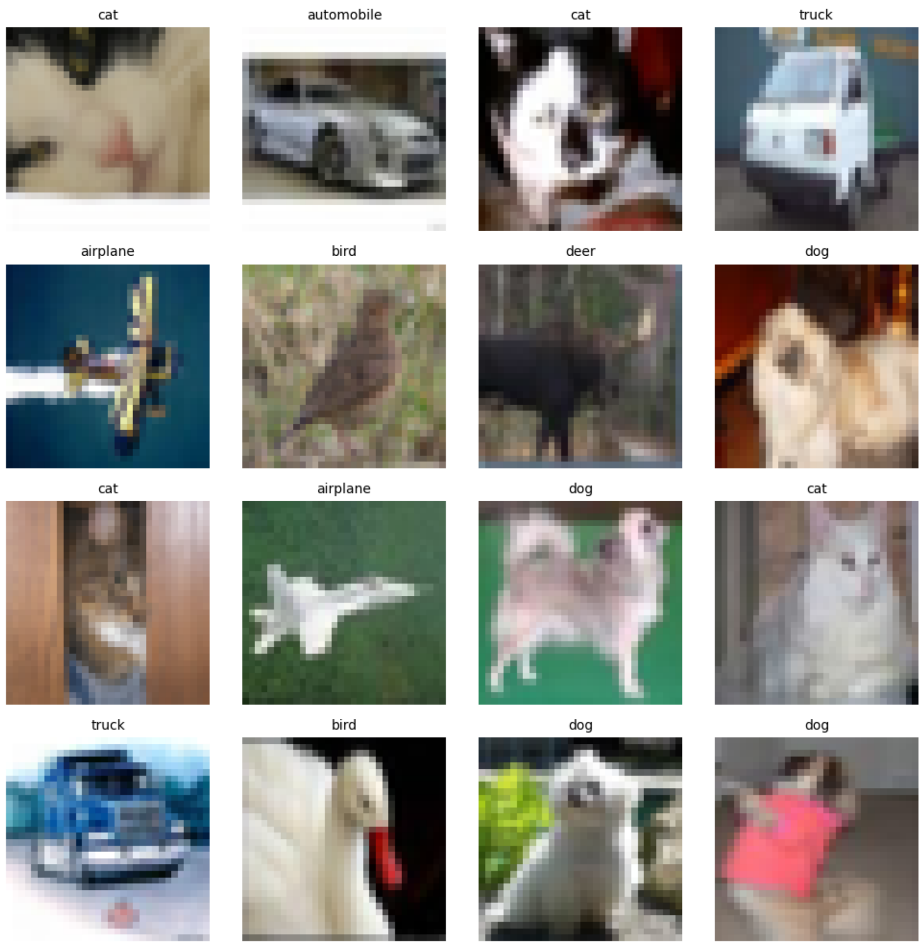

Fig. 2. Sample of CIFAR-10 dataset

`CIFAR-10` was chosen for its balanced class ratio, moderate image complexity, and appropriateness for testing lightweight and hybrid models in constrained resources, as those provided by simulated quantum circuits. Moreover, its widespread use in benchmarking DL models allows for reliable performance comparisons with existing research.

4.4. Quantum layer

The quantum component of the hybrid architecture is implemented using a parameterized quantum circuit composed of four qubits. Each qubit receives one scalar value as input, resulting in a fixed input size of four. The circuit design follows a standard pattern for hybrid applications: each qubit is first initialized with a *Hadamard gate*, which introduces superposition, and then passed through a trainable *Ry rotation gate*, allowing the circuit to learn rotational parameters during training.

This particular combination of quantum *Hadamard gates* and then *Ry* was motivated by the findings of our previous work [12], and was found to be a good trade-off in expressiveness and learning stability for small-scale quantum learning layers.

The results of the quantum circuit are the expected values as measured for each qubit. These values are aggregated into a vector and passed to the subsequent classical classification stage.

It is important to note that the quantum layer inputs are adjusted to be normalized and π-scaled, as the spin-gates tend to expect so. In this way, the classical outputs are mapped to the acceptable parameters of the quantum gates.

4.5. Classical backbone

In this paper, we compare two different classical backbone structures in the same hybrid: `LCNet050` and `MobileNetV3Small050`. Both are efficient CNNs, have been designed to be lightweight in terms of computation and small parameter size, and are thus potential candidates for realizing on quantum hardware, especially when under simulation constraints. This comparison aims to evaluate how architectural choices in classical networks influence the effectiveness of hybrid quantum-classical models. Understanding these differences is essential for optimizing performance in environments where computational resources are limited.

On the model release side, the current notebook is using `LCNet050` as the classical backbone. This architecture is composed of several convolutional layers to capture high-level abstraction from input images of size 32×32×3 (`CIFAR-10`). After the convolution feature map is extracted, the convolution feature map is flattened into a one-dimensional vector, and the linear layer to reduce its dimension is introduced to 4, which is the number of input qubits of the quantum layer. This reduction ensures compatibility with the quantum processing block and facilitates seamless integration within the hybrid architecture. The combination of compact feature representation and limited qubit input aligns with practical constraints of near-term quantum devices.

All intermediate layers in the backbone use *ReLU* activation, whereas the last linear layer ahead of the quantum circuit uses a *tanh* activation function to bound the included values in $(-1, 1)$. This is crucial because the quantum layer will interpret its inputs as angles of rotations, which are only allowed on a bounded interval.

After the quantum transform, we map the quantum output vector to ten class logits using a final fully connected layer and then use a softmax to get class probabilities. Leveraging this architecture, the experiment decouples the classical backbone's effect on model performance and conducts a fair comparison between `LCNet050` and `MobileNetV3Small050` to exclude confounding factors among different quantum settings.

4.6. Workflow

The experimental pipeline developed in this work includes different steps and is aimed at a comprehensive evaluation of hybrid quantum-classical models. The pipeline comprises the preprocessing of datasets, conditioning of the quantum layer, the training of the model, and its evaluation throughout the 5 stratified folds Cross-Validation (CV) with an Out-of-Fold (OoF) prediction strategy [13].

The main purpose of the study was to compare how the performance of such a hybrid architecture is affected by various classic backbones (`MobileNetV3Small050` and `LCNet050`). A stratified 5-fold CV method was performed in order to have a statistically robust evaluation and prevent overfitting.

In this method, the dataset was split into five non-overlapping and equal-sized datasets (folds), each fold kept the distribution of classes of the original dataset. For each of the five runs, one fold was the validation set, and the remaining four folds were used to train the model. This approach ensures that every data point is used exactly once for validation and four times for training. In addition, stratification ensures that the model is tested on all class labels fairly across folds, and this is very important to get a balanced performance measure for the multiclass classification tasks.

During each fold, predictions were recorded for the validation subset–data that the model had not encountered during training – which were then aggregated across all five iterations to construct a complete set of OoF predictions. This strategy enables the generation of model outputs for the entire dataset under strictly unseen conditions, effectively simulating real-world generalization. Each data sample contributed to the final performance evaluation exactly once, and only in the fold where it was excluded from training. The consolidated OoF prediction set was subsequently used to compute final evaluation metrics, including overall accuracy, *macro-* and *micro-averaged AUC* scores, *per-class AUC* values, and *confusion matrices*. This methodology offers a robust and unbiased estimate of model performance, free from information leakage or overly optimistic validation scores. Moreover, it provides a reliable basis for model comparison by ensuring that all performance metrics are derived from data not seen during training.

5. Results of hybrid neural network performance evaluation

In this experiment, the effectiveness of two lightweight convolutional neural networks – `MobileNetV3Small050` and `LCNet050` – was compared as classical backbones within a hybrid quantum-classical architecture using a quanvolutional preprocessing layer. Both models were trained on the `CIFAR-10` dataset with identical training configurations: 10 ***epochs***, ***batch size*** = 64, and 4-qubit quanvolutional transformation with ***stride*** = 1 and ***Receptive Field*** = 2. *AUC (macro and micro), confusion matrices*, and *class-wise AUC* scores were used as evaluation metrics.

The training pipeline involved applying a quantum preprocessing layer that outputs 5-channel quanvolutionally transformed data. This dataset was then fed into each respective backbone for classification.

To assess the influence of the backbone architecture in hybrid quantum-classical neural networks, we conducted a comparative evaluation of `MobileNetV3Small050` and `LCNet050` under a consistent quanvolutional pipeline. Table 1 summarizes the general classification performance across the `CIFAR-10` dataset using 5-fold stratified CV.

Table 1. Comparison of the general metrics after 5-fold stratified CVs

Backbone	**AUC_macro_mean**	**AUC_macro_std**	**AUC_micro_mean**	**AUC_micro_std**
MobileNetV3Small050	0,933998	0,029202	0,938082	0,001131
LCNet050	0,954004	0,024682	0,958469	0,000409
ΔAUC	0,020006		0,020387	

As shown in Table 1, when comparing hybrid quantum-classical models with various classical backbones, the `LCNet050` curve significantly dominates the `MobileNetV3Small050` across all metrics of interest. In particular, the LCNet-based model reached a higher *macro-averaged AUC* (0.9540) than its `MobileNetV3Small050` counterpart (0.9340), with an absolute difference of ΔAUC_macro = 0.02. Nevertheless, such negative improvement seems to be within the variability range defined by the*macro* AUC standard deviation of the model using `MobileNet` (AUC_macro_std = 0.0292), so the macro improvement, even present, might not be significant.

In contrast, the *micro-averaged* AUC shows a more decisive outcome: `LCNet050` recorded an AUC_micro_mean of 0.9585 versus 0.9381 for `MobileNetV3Small050`, with a corresponding ΔAUC_micro = 0.0204. Importantly, this improvement exceeds the *micro* AUC standard deviation of both models (0.0004 and 0.0011, respectively), indicating a more stable and reliable performance gain at the sample level.

The mean and standard deviation of the AUC scores over all ten `CIFAR-10` classes are given in Table 2 after 5-fold stratified CV.

Table 2. Comparison of the metrics per class after 5-fold stratified CV

Class	**AUC_M_mean**	**AUC_M_std**	**AUC_L_mean**	**AUC_L_std**	**ΔAUC_mean**
0	0,946172	0,002621	0,964261	0,001295	0,018089
1	0,967665	0,002112	0,980485	0,001488	0,012820
2	0,897933	0,002987	0,923923	0,003442	0,025990
3	0,879870	0,003302	0,905776	0,004025	0,025906
4	0,912916	0,003781	0,943832	0,001716	0,030916
5	0,914256	0,003115	0,931433	0,003443	0,017177
6	0,955699	0,001997	0,973300	0,001343	0,017601
7	0,940685	0,001617	0,966152	0,001009	0,025467
8	0,967525	0,001396	0,976660	0,001053	0,009135
9	0,957264	0,002159	0,974215	0,000966	0,016951

Across all classes, `LCNet050` consistently outperformed `MobileNetV3Small050` in terms of mean AUC, with ΔAUC_mean values ranging from $0,0091$ (class 8) to $0,0309$ (class 4). Notably, the largest improvements were observed in classes with moderately challenging features, such as $0,0309$ (class 4) and $0,0260$ (class 2), indicating `LCNet050`'s superior ability to generalize under diverse intra-class variation. These differences exceed the corresponding standard deviations of the `MobileNet` model (AUC_M_std) for these classes, suggesting that the observed improvements (Fig. 3), are not only substantial but likely statistically significant.

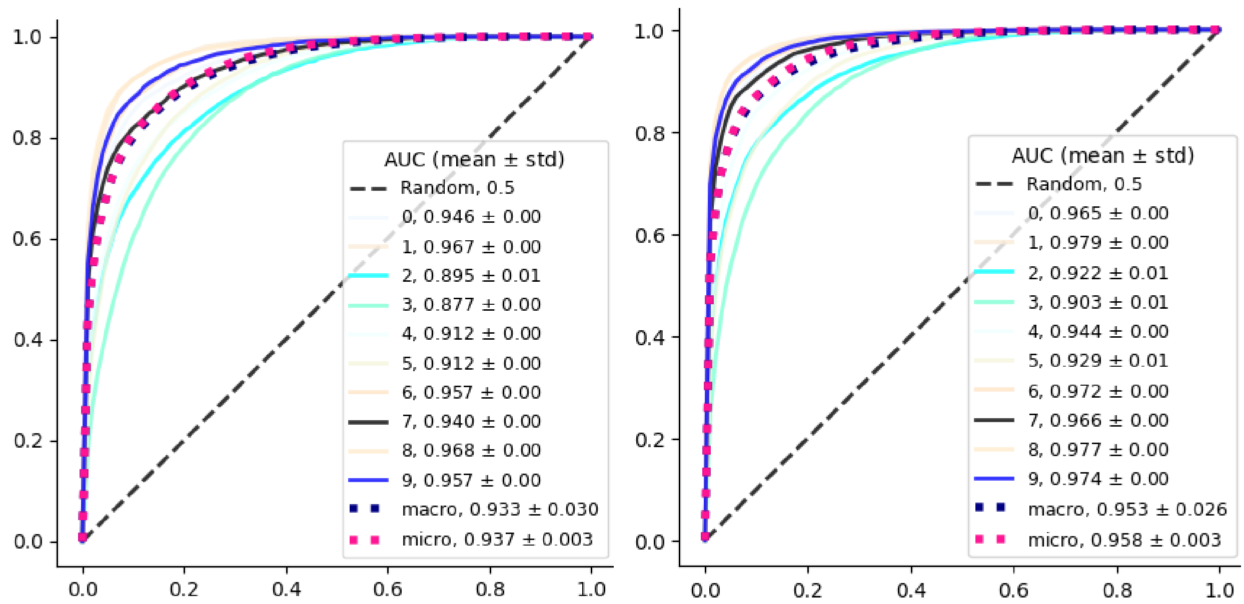

a *b*

Fig. 3. Models AUC metric: a – `MobileNetV3Small050`; b – `LCNet050`

Further, for 9 of 10 classes, `LCNet050` achieved lower standard deviation in AUC across folds as indicated by negative ΔAUC_std values. This implies that `LCNet050` not only generates better quality video classification on average, but also provides predictions that are more stable across different test splits – an important consideration for desiring consistency across training splits. The maximum decrease of $AUCs$ spread was also obtained for class 4 (-0.00207), confirming again its good versatility in presence of more ambiguous samples. In addition to the CVmetrics, we assessed the generalization performance of both pre-trained architectures via OoF predictions. The *macro/micro-averaged* AUC scores based on the OoF validation sets during 5-fold training are reported in Table 3.

Table 3. Comparison of the general metrics after OoF training and validation

Class	AUC_M_OoF	AUC_M_std (from CV)	AUC_L_OoF	AUC_L_std (from CV)	ΔAUC_OoF
0	0,953404	0,002621	0,966818	0,001295	0,013414
1	0,96999	0,002112	0,984556	0,001488	0,014566
2	0,908798	0,002987	0,928879	0,003442	0,020081
3	0,895899	0,003302	0,916242	0,004025	0,020343
4	0,923202	0,003781	0,951818	0,001716	0,028616
5	0,920091	0,003115	0,93697	0,003443	0,016879
6	0,96203	0,001997	0,971545	0,001343	0,009515
7	0,951162	0,001617	0,970626	0,001009	0,019464
8	0,972111	0,001396	0,979384	0,001053	0,007273
9	0,964333	0,002159	0,978232	0,000966	0,013899

The `LCNet050` backbone takes clear advantage again to achieve the AUC_macro of 0.9585 and AUC_micro of 0.9634, against 0.9421 and 0.9465 for the `MobileNet`-based model. These gap values are ΔAUC_macro = 0.0164 and ΔAUC_micro = 0.0168, thus supporting again the consistent higher performance of `LCNet` in class-wise and average prediction resolution too. The macro level improvement does not outweigh the macro $AUCs$ standard deviation for `MobileNetV3Small050` (0.0262), but it is instructive that the micro improvement is significant beyond the observed variability, i.e., ΔAUC_micro surpasses the AUC_micro_std of these models (0.001013 and 0.000380, respectively). This indicates meaningful and consistent enhancement of the performance for the prediction of the `LCNet`-based model when operating over the whole logical space. The generally small *stds* of the `LCNet` results further confirm its robustness over OoF folds.

Table 4. Comparison of the metrics per class after OoF training and validation

Class	AUC_M_OoF	AUC_M_std (from CV)	AUC_L_OoF	AUC_L_std (from CV)	ΔAUC_OoF
0	0,953404	0,002621	0,966818	0,001295	0,013414
1	0,96999	0,002112	0,984556	0,001488	0,014566
2	0,908798	0,002987	0,928879	0,003442	0,020081
3	0,895899	0,003302	0,916242	0,004025	0,020343
4	0,923202	0,003781	0,951818	0,001716	0,028616
5	0,920091	0,003115	0,93697	0,003443	0,016879
6	0,96203	0,001997	0,971545	0,001343	0,009515
7	0,951162	0,001617	0,970626	0,001009	0,019464
8	0,972111	0,001396	0,979384	0,001053	0,007273
9	0,964333	0,002159	0,978232	0,000966	0,013899

The per-class AUC obtained via OoF predictions is reported in Table 4. Remarkably, the

`LCNet050` backbone achieves higher AUC scores across all ten `CIFAR-10` classes, with ΔAUC_OoF improvements ranging from 0.0073 (class 8) to 0.0286 (class 4). Crucially, in every class, the performance gain exceeds the standard deviation (AUC_M_std) recorded for `MobileNetV3Small050` during CV, indicating that the improvements are not only consistent but also statistically significant. The most pronounced gains are observed in classes 2 (frog), 3 (cat), and 4 (deer), with improvements of 0.0201, 0.0203, and 0.0286, respectively–highlighting `LCNet050`'s ability to better generalize in moderately challenging categories. Even in already well-performing classes, such as class 1 (automobile), class 8 (ship), and class 9 (truck), the `LCNet050`-based hybrid maintains measurable advantages. This consistent outperformance across the full class spectrum strongly suggests that the `LCNet` backbone achieves not only higher accuracy but also more reliable per-class discrimination, making it better suited to extract and interpret quantum-encoded features.

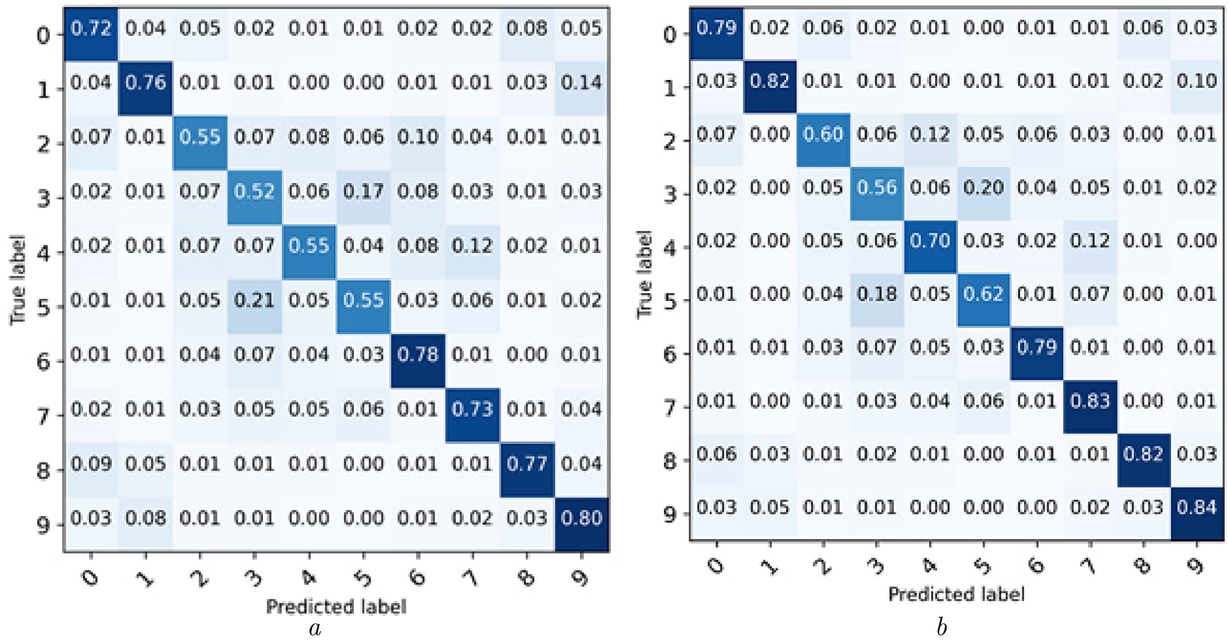

a *b*

Fig. 4. Models normalized confusion matrices: *a* – `MobileNetV3Small050`; *b* – `LCNet050`

Confusion matrices (Fig. 4), also confirm the enhanced classification accuracy of the `LCNet050`-based QNN model, with a sharper diagonal and reduced inter-class confusion.

6. Discussion of the results of hybrid neural network evaluation

6.1. Accuracy and generalization analysis

We experimentally validate that the traditional backbone structure makes a significant difference in the performance of hybrid quantum-classical neural models. For all CV and OoF metrics, the model based on `LCNet050` reported consistently higher *macro-* and *micro-AUC* values than the one of `MobileNetV3Small050`. Notably, `LCNet050` achieved a gain of around 0,02 (both *macro AUC* and *micro AUC*) on average over the evaluated settings.

In addition, `LCNet050` achieved higher per-class $AUCs$ across all ten `CIFAR-10` classes, demonstrating consistent superiority over `MobileNetV3Small050`. The most notable improvements were observed in mid-performing classes such as class 2 (frog) and class 4 (deer), where the performance gains were particularly pronounced. These results highlight `LCNet050`'s strong generalization capabilities, not only for dominant and easily separable classes but also for those with higher intra-class variability. This consistent performance across the entire class distribution reflects a more balanced and robust classification strategy, reinforcing `LCNet050`'s suitability as a backbone for hybrid quantum-classical models.

The confusion matrix analysis is also in line with this view, showing the decrease of inter-class confusion and the increased distribution of diagonal elements.

6.2. Performance stability across folds

Aside from its better mean performance, the `LCNet050` backbone also showed far greater stability across both CV folds and individual classes. The standard deviations of *macro*- and *micro-AUC* scores were consistently lower for `LCNet050` under both CV and OoF evaluations, indicating its stability. Notably, 9 out of 10 classes showed variance reduction in AUC with the use of `LCNet050`, which further enhances its capacity to deliver consistent performance under different training splits. The greatest decrease in variability was experienced by class 4, where the standard deviation decreased by a value of approximately 0.00207, once more demonstrating the stability of the model in its ability to handle more complex or open-ended classes. These decreased predictive variances coupled with the continued gains in accuracy indicate that `LCNet050` not only makes more accurate predictions but also greater robustness and reliability–qualities paramount to applying hybrid models in realistic, practical applications.

6.3. Classical-quantum integration compatibility

These results indicate that classical architecture is not a trivial choice when considering the design of hybrid quantum-classical models, but something that must be carefully thought of. Small variations in the backbone design can still have an impact on accuracy, stability, and per-class sensitivity. With the increasing access to quantum devices, such design choices will be of utmost importance to establish the practicability and success of hybrid models in production.

Next steps could involve investigating different types of quantum layers (e.g., variational circuits and entanglement schemes) interacting with classical backbone properties due to depth, connectivity, and *receptive field* design. This will enable a more holistic exploration of the design space for hybrid architectures and allow the development of principled rules for the construction of high-performance models.

Conclusions

Based on the results of the study, the following conclusions can be drawn:

1. The classical backbone architecture significantly influences the performance of hybrid quantum-classical neural networks. Despite identical quantum preprocessing layers, the `LCNet050` backbone consistently outperformed `MobileNetV3Small050` across all evaluated metrics. This highlights that classical design decisions remain a critical factor in hybrid model effectiveness.

2. `LCNet050` offers both higher accuracy and greater stability. The `LCNet050`-based hybrid model not only achieved superior *macro*- and *micro-AUC* scores but also demonstrated reduced variance across folds and classes, indicating more reliable generalization.

3. Careful integration between quantum and classical components is essential. The observed performance gains suggest that architectural compatibility–particularly how well the classical model processes quantum-transformed inputs plays a key role in hybrid network success. These findings emphasize the importance of backbone selection in hybrid architectures and provide a foundation for future studies exploring broader architectural combinations and quantum configurations [15].

Acknowledgements

This research was supported by the National Research Fund of Ukraine under grant 2025.06/0100. The grant covered the development of neural network models for classification and object detection. It was also supported by the Ministry of Education and Sciences of Ukraine (MESU), grant 2715r. This part focused on exploring advanced deep learning methods based on hardware-specific requirements.

References

[1] H. I. G. Hernández, R. T. Ruiz, and G. Sun, "Image Classification via Quantum Machine Learning," *arXiv preprint*, 2020. [Online]. Available: https://doi.org/10.1155/2015/918045

[2] S. An, C. Lee, H. Moon, and J. Park, "An Ensemble of Simple Convolutional Neural Network Models for MNIST Digit Recognition," *arXiv preprint*, 2020. https://doi.org/10.48550/arxiv.2008.10400

[3] S. Resch and U. R. Karpuzcu, "Quantum Computing: An Overview Across the System Stack," *arXiv preprint*, 2019. https://doi.org/10.48550/arxiv.1905.07240

[4] T. Begušić and G. K.-L. Chan, "Fast classical simulation of evidence for the utility of quantum computing before fault tolerance," *arXiv preprint*, 2023. https://doi.org/10.48550/arxiv.2306.16372

[5] A. Mari, T. Bromley, J. Izaac, M. Schuld, and N. Killoran, "Transfer Learning in Hybrid Classical-Quantum Neural Networks," *Quantum*, vol. 4, p. 340, Oct. 2020. https://doi.org/10.22331/q-2020-10-09-340

[6] Z. Chen, R. Zhang, L. Yu, and Y. Gao, "VQNet: Library for a Quantum-Classical Hybrid Neural Network," *arXiv preprint*, 2019. https://doi.org/10.48550/arxiv.1901.09133

[7] K. He, X. Zhang, S. Ren, and J. Sun, "Deep Residual Learning for Image Recognition," arXiv preprint, 2015. https://doi.org/10.48550/arxiv.1512.03385

[8] C. Cui, J. Wang, B. Wu, and Y. Gao, "PP-LCNet: A Lightweight CPU Convolutional Neural Network," *arXiv preprint*, 2021. https://doi.org/10.48550/arxiv.2109.15099

[9] D. Lykov, R. Schutski, A. Galda, V. Vinokur and Y. Alexeev, "Tensor Network Quantum Simulator With Step-Dependent Parallelization," *arXiv (Cornell University)*, 2020. https://doi.org/10.48550/arxiv.2012.02430

[10] A. De Lorenzis, A. Giusti, M. Pierini, and S. Carrazza, "Harnessing quantum extreme learning machines for image classification," *Phys. Rev. Applied*, vol. 23, no. 4, Apr. 2025.https://doi.org/10.1103/PhysRevApplied.23.044024

[11] D. Shepherd, "On the Role of Hadamard Gates in Quantum Circuits." *In arXiv:quant-ph/0508153*, Mar. 2006.https://doi.org/10.48550/arXiv.quant-ph/0508153

[12] Y. Gordienko, A. Khmelnytskyi, V. Taran, and S. Stirenko, "Hybrid Neural Networks with Multi-channel Quanvolutional Input for Medical Image Classification," *in Trends in Sustainable Computing and Machine Intelligence*, ICTSM 2024. Algorithms for Intelligent Systems. Singapore: Springer, pp. 181–193, 2025. https://doi.org/10.1007/978-981-96-1452-3_15

[13] Kakkar, Sushma, A. Jaya Lakshmi, Anand Deva Durai C, Wakeel Ahmad, Mohamed Uvaze Ahamed Ayoobkhan, P. Veeramanikandan, P. Nagasekhar Reddy, Vijay Kumar Dwivedi, and A. Rajaram, "Enhancing Energy Efficiency and Classification Modeling Through a Combined Approach of LightGBM and Stratified KFold Cross-Validation," *Electric Power Components and Systems*, pp. 1–19, 2024. https://doi.org/10.1080/15325008.2024.2315213

[14] IBM Quantum, "Qiskit Documentation." [Online]. Available: https://qiskit.org/documentation/

[15] "Kaggle notebook for MobileNetV3Small050." [Online]. Available: https://www.kaggle.com/code/arseniykhmelnitskiy/cifar10-full-q5-w4-mobnv3-mb1-batch-64-t

УДК 004.415.2

ПОРІВНЯЛЬНИЙ АНАЛІЗ АРХІТЕКТУР LCNET050 ТА MOBILENETV3 У ГІБРИДНИХ КВАНТОВО-КЛАСИЧНИХ НЕЙРОННИХ МЕРЕЖАХ ДЛЯ КЛАСИФІКАЦІЇ ЗОБРАЖЕНЬ

Арсеній Хмельницький
https://orcid.org/0009-0009-1631-8172

Юрій Гордієнко
https://orcid.org/0000-0003-2682-4668

Національний технічний університет України
«Київський політехнічний інститут імені Ігоря Сікорського», Київ, Україна

У даному дослідженні розглянуто вплив класичної архітектури базової моделі на продуктивність гібридних квантово-класичних нейронних мереж) у задачах класифікації зображень. Гібридні моделі поєднують представницьку здатність класичного глибокого навчання з потенційними перевагами квантових обчислень. Зокрема, використано архітектуру квантово-конволюційної нейронної мережі, де квантовий конволюційний шар на основі чотирикубітного Ry-ланцюга виконує попередню обробку зображень перед класичними обчисленнями. Попри зростаючий інтерес до гібридних моделей, досі бракує систематичних досліджень, які б вивчали вплив варіацій у класичній архітектурі на загальну ефективність гібридних квантово-класичних нейронних мереж. Щоб заповнити цю прогалину, ми порівнюємо дві легкі конволюційні архітектури – `MobileNetV3Small050` та `LCNet050` – інтегровані з однаковим квантовим шаром попередньої обробки. Обидві моделі оцінюються на наборі даних `CIFAR-10` з використанням 5-кратної стратифікованої крос-валідації. Ефективність вимірюється за допомогою кількох метрик, включно з точністю, макро- та мікросереднім показником площі під кривою, а також матрицями невідповідностей по класах. Результати демонструють, що гібридна модель на базі `LCNet` стабільно перевершує `MobileNet` за загальною точністю та значеннями площі під кривою, забезпечуючи кращий баланс між класами та стійкість у розпізнаванні менш представлених класів. Це підкреслює критичну роль вибору класичної архітектури в гібридних квантово-класичних системах. Незважаючи на фіксований квантовий шар, взаємодія між квантовою попередньою обробкою та класичним вилученням ознак суттєво впливає на результати моделі. Це дослідження робить внесок у зростаючий корпус робіт у сфері квантово-підсилених навчальних систем, демонструючи важливість вибору класичної архітектури. Майбутні дослідження можуть розширити ці результати на інші набори даних, глибші або трансформерні архітектури, а також більш виразні квантові схеми.

Ключові слова: нейронні мережі, квантові обчислення, гібридні нейронні мережі, класифікація зображень.