

# Understanding the effects of data encoding on quantum-classical convolutional neural networks

Maureen Monnet\*, Nermine Chaabani\*, Theodora-Augustina Drăgan\*, Balthasar Schachtner†, Jeanette Miriam Lorenz\*

\*Fraunhofer Institute for Cognitive Systems IKS, Munich, Germany

†LMU University Hopsital, Munich, Germany

{maureen.monnet, theodora-augustina, jeanette.miriam.lorenz}@iks.fraunhofer.de  
balthasar.schachtner@med.lmu.de

**Abstract**—Quantum machine learning was recently applied to various applications and leads to results that are comparable or, in certain instances, superior to classical methods, in particular when few training data is available. These results warrant a more in-depth examination of when and why improvements can be observed. A key component of quantum-enhanced methods is the data encoding strategy used to embed the classical data into quantum states. However, a clear consensus on the selection of a fitting encoding strategy given a specific use-case has not yet been reached. This work investigates how the data encoding impacts the performance of a quantum-classical convolutional neural network (QCCNN) on two medical imaging datasets. In the pursuit of understanding why one encoding method outperforms another, two directions are explored. Potential correlations between the performance of the quantum-classical architecture and various quantum metrics are first examined. Next, the Fourier series decomposition of the quantum circuits is analyzed, as variational quantum circuits generate Fourier-type sums. We find that while quantum metrics offer limited insights into this problem, the Fourier coefficients analysis appears to provide better clues to understand the effects of data encoding on QCCNNs.

**Index Terms**—quantum machine learning, data encoding, quantum convolutional neural networks, medical imaging

## I. INTRODUCTION

Quantum convolutional neural networks (QCNNs) are a promising architecture in the domain of image classification that is theoretically shown to be beneficial in situations where few training data points are available [1]. This makes them particularly interesting for e.g. medical applications where data availability is limited due to data collection and annotation costs, as well as privacy and security concerns. While QCNNs are originally designed to be fully implemented on quantum computers [2], the currently available quantum hardware – commonly referred to as Noisy Intermediate-Scale Quantum (NISQ) computers – are not yet up to the task. More technological development is required in order to implement robust fully-quantum solutions, such as an increased number of qubits, reduced error rates and improved coherence times.

A more NISQ-friendly approach to QCNNs is achieved by running a hybrid quantum-classical version, where one or several chosen components of the classical convolution neural network (CNN) architecture are replaced by a quantum variant. In the classical case, a CNN typically comprises convolutional layers responsible for extracting characteristic features from

images, followed by fully connected layers employed for the classification. In the quantum-classical case, the convolutional layers stand out as promising candidates for their mapping onto a NISQ quantum computer, as they handle data sequentially in smaller segments, in contrast to the linear layers that process all features simultaneously. In that scenario, the convolutional layer is replaced by a variational quantum circuit (VQC) containing trainable operations, while the remainder of the architecture remains classical. The optimization is then performed on a classical computer using a classical optimizer.

In order to process classical data with a quantum circuit, a crucial step is the data encoding, as it defines the feature space the quantum circuit is able to generate ([3], [4]). The choice of a suitable encoding strategy for a given dataset is therefore a non-trivial endeavour, with literature displaying ongoing efforts to better understand its impact on the performance of the entire algorithm [5]. Further techniques related to data encoding, such as data-reuploading, have demonstrated the potential to enhance performance significantly [6]. In this approach, the encoding scheme is repeated between the different variational layers.

In this work, we evaluate the performance on two medical imaging datasets of QCCNNs with four different encoding strategies, namely the angle encoding around the X- and Y-axis, the higher-order, and the amplitude encoding. The studied datasets include 2D ultrasound images of the breast for identifying malignant lesions, alongside abdominal computed tomography (CT) scans for conducting multi-class classification of body organs. Specifically, we compare the performance of the selected data encoding strategies when they are re-uploaded up to five times. We then strive to understand the differences in performances by employing two distinct approaches: in the first approach, we use quantum metrics proposed in literature to quantify the power of quantum neural networks. In the second approach, we apply the Fourier formalism to the outputs generated by our quantum circuits. The motivation behind pursuing the latter arises from the observation that outputs generated by variational quantum circuits can be expressed in terms of Fourier series [7].

Our contributions are thus as follows:

- We analyze the effect of four different data encodings on the performance of a quantum-classical convolutional

neural network (QCCNN) for binary and multi-class classifications of ultrasound and CT medical images. We additionally investigate the impact of input scaling and data-reuploading on the results.

- We analyze the performance of all architectures in the light of 3 relevant quantum metrics, namely the expressibility, the entanglement capability and the normalized effective dimension to find possible correlations.
- We extend this analysis by applying the Fourier formalism on the outputs produced by the quantum circuits.

The paper is structured as follows. We present the related work in section II. The background information relevant to the designed QCCNNs and their encodings, the data-reuploading scheme, the three quantum metrics and the Fourier series representation of quantum circuits is introduced in section III. We present the datasets used, the performances of the QCCNNs and their potential correlations with the defined quantum metrics and Fourier coefficients in the results and discussion section IV. We conclude in section V.

## II. RELATED WORK

### A. Empirical studies on the performance of QCCNNs

Fully quantum QCCNNs have been studied in several works ([2], [8], [9], [10], [11], [12]). However, with the current limitations posed by NISQ hardware, it is of interest to move only specific components of the architecture to the quantum computer, resulting in a hybrid quantum-classical architecture. The convolutional layer in the classical CNN is substituted with either an untrainable or trainable quantum convolutional layer in ([13], [5], [14], [15], [16]), where good performance is shown on the MNIST dataset in ([13], [5]). [5] particularly focuses on the impact of the encoding scheme and filter size of the quantum convolution, showing that some architectures perform better than others. Building upon this, our work in [15] and [16] introduces diverse VQC architectures with varied encoding schemes across multiple medical datasets. We notably show that the choice of ansatz significantly impacts the performance of the algorithm, and start to examine some theoretical clues to elucidate why this is the case.

### B. Performance analysis with theoretical clues

Different metrics have been proposed in literature to quantify certain properties of quantum circuits, such as the expressibility, the entanglement capability and the effective dimension ([17], [18]). Work to link these properties to the performance of the VQC has not yet shown conclusive correlation. In [19], the authors study the influence of the chosen ansatz on the prediction quality and speed of convergence of a hybrid quantum reinforcement learning algorithm. We studied in [16] the link between various quantum pooling strategies to reduce the dimension of the data and the effective dimension of the circuit, again without finding a correlation between the two. This still raises the question as to whether these metrics might be relevant for choosing a good embedding strategy.

Alternatively, it is possible to look at the problem from a different angle: previous work in ([7], [20]) shows that QML




Fig. 1: The circuit structure used in this study.

models that use the variational approach generate Fourier-type sums. From a signal processing point of view, this means that the outputs of quantum circuits can be decomposed into their constituting sine and cosine waves, each with a weighted contribution to the overall output. In [21], the authors analyze this formalism and its implications for different types of circuit architectures commonly used in QML for both one-dimensional and multi-dimensional inputs, while using a varying number of data encoding and ansatz layers. The authors of [22] compare the Fourier coefficient distributions for a 2-qubit circuit and a similar one with an additional ancillary qubit, revealing distinct distributions. This suggests that they learn different function classes despite only varying in qubit count. They also observe that measurements on different qubits in the 2-qubit circuit yield varying coefficient distributions.

To the best of our knowledge, our study represents the first attempt to analyze the performance of a QCCNN using diverse encoding strategies, considering both the quantum metrics mentioned earlier and Fourier decomposition.

## III. BACKGROUND

In this section, we provide the relevant information for the construction of the QCCNNs, specifically addressing the encodings employed and the data re-uploading strategy. Following that, we define the quantum metrics used in this work, and describe the Fourier representation of quantum circuits.

### A. QCCNNs and encodings within this work

We build on the study conducted in [15] and construct a QCCNN by replacing the first classical convolutional layer of a CNN by a VQC, while keeping the remainder of the architecture unaltered. All qubits are then measured and their outcomes are mapped onto separate output channels. A classical fully-connected layer performs the final classification using the sigmoid activation function.

Fig. 1 shows the general structure of this circuit. The VQCs in this study are designed as follows: we first define an encoding strategy  $S(x)$ , then we select an ansatz  $W(\theta)$  to remain unchanged for every encoding tested, as we are

interested in observing the effect of the data embedding. We choose a basic entangling layer as it combines simplicity with high performance ([15], [16]). In this ansatz, each qubit  $i$  is rotated by a trainable angle  $\theta_i$  around the X-axis before applying a sequence of CNOT entangling gates on adjacent qubit pairs. Note that for a 2x2 filter size in the convolutional layer, the number of qubits required to construct such a filter may vary depending on the choice of encoding, e.g. 4 qubits for the angle and higher-order encoding, versus 2 for the amplitude encoding. For all encoding strategies used in this study, a rescaling factor  $f$  is multiplied with the input data.

The following encoding techniques are explored:

- 1) **Angle Encoding:** This embedding method is popular particularly due to its simple implementation and typically satisfactory performance [15]. For an N-entry classical data vector  $\mathbf{x} = (x_1, \dots, x_N)$  we can encode the data as

$$|\Psi\rangle = \bigotimes_{i=1}^N \cos(x_i) |0\rangle + \sin(x_i) |1\rangle \quad (1)$$

and we can write the state preparation unitary as  $\mathcal{U} = \bigotimes_{i=1}^N \exp\left(\frac{-i\phi\sigma_j}{2}\right)$  with  $\sigma_j$  the Pauli matrices, generators of the Pauli gates. While the execution of this encoding is relatively straightforward, there is no efficiency benefit in the number of required qubits, with an N-dimensional input requiring an N-qubit system. In this work, we use two possible angle encoding schemes, namely with an  $R_X$  gate or an  $R_Y$  gate, corresponding to  $\sigma_X$  and  $\sigma_Y$  respectively.

- 2) **Higher-Order Encoding:** First proposed in [4], this encoding was initially designed for Support Vector Machine (SVM) classifiers where it was shown that a quantum advantage can only be achieved if the encoding feature map is classically hard to simulate. The feature map unitary as used in [18] is given by

$$|\Psi\rangle = H^{\otimes n} \exp\left(i \sum_{S \subseteq [n]} \phi_S(\vec{x}) \prod_{i \in S} Z_i\right). \quad (2)$$

A Hadamard gate is first applied on each qubit, followed by a first-order angle encoding using  $R_Z(\phi_1(x_i))$  gates with  $\phi_1(x_i) = f * x_i$  and a pairwise second-order operation  $R_{ZZ}(\phi_2(x_i, x_j))$  with  $\phi_2(x_i, x_j) = f * x_i x_j$ .

- 3) **Amplitude Encoding:** Classical data can be represented as amplitudes of a quantum state. For a classical data vector  $x = (x_1, \dots, x_N)$ , this can be written as

$$|\Psi\rangle = \sum_{i=1}^N x_i |i\rangle \quad (3)$$

where  $|\Psi\rangle$  is the resulting quantum state,  $x_i$  are the amplitudes containing the classical information such that  $\sum_{i=1}^N |x_i|^2 = 1$ , and  $|i\rangle$  are the quantum states for each vector component, with  $i$  being the index written in

binary. An N-dimensional amplitude encoded classical vector requires  $\log_2(N)$  qubits. This implies that we only require 2 qubits to encode an input of size 4. This reduction in the required number of qubits is the main advantage of amplitude encoding. To prepare the amplitudes, a quantum state preparation is used. In this work, we use the Möttönen state preparation [23] as implemented in pennylane [24].

## B. Data-reuploading

Unlike classical neural networks, which are able to process the original data multiple times in their hidden layers, quantum neural networks are limited by the no-cloning theorem. To circumvent this limitation, [6] propose a single-qubit universal classifier consisting of an encoding and a trainable unit through which the data is repeatedly passed. For a general single-qubit rotation  $U(\theta, \phi, \delta)$  which we consider as the processing unit, and an arbitrary data encoding rotation gate  $U(x)$ , we can define a single layer as  $L(x, \theta, \phi, \delta) = U(\theta, \phi, \delta) U(x)$ . This constitutes the building block of the classifier and can be applied repeatedly. The circuit unitary for  $N$  layers can be written as  $\mathcal{U}(\mathbf{x}, \theta, \phi, \delta) = L_1(x_1, \theta_1, \phi_1, \delta_1) \dots L_N(x_N, \theta_N, \phi_N, \delta_N)$ . The principle can also be extended to a multi-qubit classifier which enables entangling operations between qubits, improving the performance and lowering the number of required layers.

## C. Quantum Metrics

Multiple metrics to describe the properties of quantum circuits were proposed in literature. In this work, we investigate three quantum metrics to identify their possible correlations with the choice of an encoding strategy.

- 1) **Expressibility** [17]: The expressibility quantifies the ability of a circuit to explore the Hilbert space. For a single qubit, a highly expressible VQC is able to explore the entire Bloch Sphere uniformly. The degree of uniformity is measured by randomly sampling parameters  $\theta$  of the VQC and taking the statistical distance measure Kullback–Leibler-divergence relative to the ensemble of Haar random states. If a circuit is highly expressible, this distance is low. This is given by

$$Expr = D_{KL} \left( \hat{P}_{VQC}(F; \theta) \| P_{Haar}(F) \right) \quad (4)$$

where the fidelities  $F$  are obtained by repeating the procedure of uniformly sampling doublets of parameter vectors and computing their corresponding state overlap.  $\hat{P}_{VQC}(F; \theta)$  is the estimated probability distribution of fidelities obtained by sampling from the VQC and  $P_{Haar}(F)$  is known and obtained analytically by  $P_{Haar}(F) = (N-1)(1-F)^{N-2}$  [25].

- 2) **Entanglement capability** [17]: This metric is given by the Meyer-Wallach (MW) entanglement measure and quantifies the capacity of a VQC to generate entangled states. It would be equal to zero for a circuit that generates product states (no entanglement) and one in

the case of an output consisting only of maximally entangled states. The MW measure is described by

$$Q(|\psi\rangle) \equiv \frac{4}{N} \sum_{j=1}^N D(\iota_j(0)|\psi\rangle, \iota_j(1)|\psi\rangle) \quad (5)$$

where  $N$  is the number of qubits in the system and  $\iota_j(b)$  is a linear mapping that acts on the state in the computational basis ( $b \in \{0, 1\}$ ) with the  $j$ -th qubit absent using the Kronecker delta as

$$\iota_j(b) |b_1 \dots b_N\rangle = \delta_{bb_j} |b_1 \dots b_{j-1} b_{j+1} \dots b_N\rangle. \quad (6)$$

The generalized distance  $D$  is computed as

$$D(|u\rangle, |v\rangle) = \frac{1}{2} \sum_{i,j} |u_i v_j - u_j v_i|^2. \quad (7)$$

The final numerical value used to estimate the entanglement capacity is computed using the following equation

$$Ent = \frac{1}{|S|} \sum_{\theta_i \in S} Q(|\psi_{\theta_i}\rangle). \quad (8)$$

where the average of the MW measure is taken over the set  $S$  of sampled circuit parameter vectors  $\theta_i$ .

- 3) **Effective Dimension (ED)** [18]: The ED is designed to evaluate the information capacity of a statistical model. When considering ML models as such, the actual dimension they explore may be a subset of the model space. The metric characterizing this space is given by the Fisher Information Matrix (FIM). A higher effective dimension can be interpreted as corresponding to a model exploring more of the space of all its possible functions. We define a  $d$ -dimensional real Riemannian parameter space  $\Theta$  for our statistical model  $\mathcal{M}_\Theta$  with parameters  $\theta \in \Theta \subset [-1, 1]^d$  and consider input-output data pairs  $(x, y)$  with  $x \in \mathcal{X}^{in}$  and  $y \in \mathcal{Y}^{out}$  which are related for a fixed choice of parameter  $\theta$  as  $p(x, y; \theta) = p(y|x; \theta)p(x)$ . One can capture the amount of information acquired from such parameterization using the FIM  $F(\theta) \in \mathbb{R}^{d \times d}$ , described by

$$F(\theta) = \mathbb{E} \left[ \frac{\partial}{\partial \theta} \log p(x, y; \theta) \frac{\partial}{\partial \theta} \log p(x, y; \theta)^\top \right] \quad (9)$$

Geometrically, the FIM captures the curvature of the log-likelihood function and it is a gauge for the sensitivity of the model to movements in the parameter space. A high FIM reflects a steeper curvature and thus greater information content. With the help of the FIM as a metric for the parameter space  $\Theta$ , we can now define the effective dimension as a bounded capacity measure for a number  $n > 1$  of data samples as

$$d_{eff}(\mathcal{M}_\Theta) := \frac{2 \log \left( \frac{1}{V_\Theta} \int_\Theta \sqrt{\det \left( \text{id}_d + \frac{\gamma n}{2\pi \log n} \hat{F}(\theta) \right)} d\theta \right)}{\log \left( \frac{\gamma n}{2\pi \log n} \right)} \quad (10)$$

In order to obtain a quantity that can be used to compare the different models, the effective dimension is then divided by the number of parameters and is referred to as normalized effective dimension (NED).

#### D. Quantum circuits and Fourier series

One way to understand QML models constructed with VQCs is through a Fourier series representation [7]. This is especially relevant for models that use data-reuploading schemes with encoding gates of the form  $S(x) = e^{ixH}$ , where  $x \in \mathbb{R}^M$  is the input and  $H$  is a general Hamiltonian for an arbitrary dimension  $M \in \mathbb{N}$  of the input data. The function describing the model can thus be written as

$$f_\theta(\mathbf{x}) = \sum_{\mathbf{w} \in \Omega} c_{\mathbf{w}}(\theta) e^{i\mathbf{w}\mathbf{x}}, \quad (11)$$

such that  $\mathbf{x}$  is a  $M$ -dimensional input vector,  $\mathbf{w}$  is a  $M$ -dimensional frequency vector in the  $\Omega \subset \mathbb{R}^N$  frequency spectrum determined by the eigenvalues of the data encoding Hamiltonians,  $\mathbf{w}\mathbf{x}$  is a scalar product and the coefficients  $c_{\mathbf{w}} \in \mathbb{C}$  are determined by the remainder of the circuit and satisfy the condition  $c_{\mathbf{w}} = c_{-\mathbf{w}}^*$ .

In order to extend the frequency spectrum of a quantum circuit and therefore to increase the degree of the Fourier series it generates, one may increase the number of qubits in the circuit, or the number of reuploading layers. In the ideal case, a chosen VQC architecture should be able to generate a rich frequency spectrum and adapt the coefficients of the series in such a way that it can approximate the target function well. However, not all architectures are able to contain enough adaptable degrees of freedom to appropriately adjust the coefficients of the series they generate. A detailed study of how the choice of the model architecture affects the number of parameters can be found in [21]. In this work, we are solely focused on what the authors refer to as a Parallel Ansatz (PA), where each data feature is encoded in a different qubit with a single-qubit gate. Following the notations used in [21], the number of parameters that the PA contains is given by

$$N_p^{PA} = (d^{2M} - 1)(L + 1) \quad (12)$$

where  $d$  is the local dimension given by  $d = 2^n$  for a qubit system in which  $n$  is the dimensionality of the local space and depends on the encoding,  $M$  is the number of qubits, and  $L$  is the number of layers. The degrees of freedom for a general Fourier series of degree  $D$  is given by

$$\nu = (2D + 1)^M \quad (13)$$

with  $D = (d - 1)L$ . The condition  $N_p^{PA} \geq \nu$  is to be met for the circuit to adapt all coefficients in the series.

#### IV. RESULTS AND DISCUSSION

We proceed to present the performances of our developed QCCNNs on two medical datasets, considering different encodings and varied scalings of the inputs. Subsequently, we analyze these performances considering three quantum metrics

typically employed in literature, and using the Fourier formalism, which describes the function class that quantum circuits give rise to. All experiments in this work were performed in simulation using the PennyLane software library [24] in conjunction with PyTorch [26] in a noiseless environment.

#### A. Datasets

In this study, we work with relatively small datasets. This requirement comes both from the application side, as medical imaging datasets usually are in the order of 100 to 1000 images, and from the method side, since quantum variants of CNNs were shown to generalize well even when little training data is available. We select two datasets from the MedMNIST benchmark ([27], [28]).

The first dataset is BreastMNIST, which contains 546 training and 78 validation images of 2D ultrasounds of the breast in a 28x28 pixels resolution. The images display breast tissues that are either normal, malignant or benign. The task is to perform binary classification, where the benign and normal lesions are combined into one class.

Our second dataset is OrganAMNIST, which consists of cropped and resized images to a 28x28 resolution based on 3D abdomen CT scans from the Liver Tumor Segmentation Benchmark (LiTS) [29]. A subset of 1,000 training and 600 validation images is taken for this study. The task is to perform multi-class classification of 11 organs.

#### B. Performances of the QCCNNs

1) *Performance with varying input scalings:* The scaling of the classical input data before encoding is not a trivial choice and impacts both the function expressed by the circuit [3] and the circuit's coverage of the Hilbert space. In this work, the pixel values of the images are normalized between -1 and 1 before a scaling factor  $f$  is applied. Due to the excessively long training times, the study of the performance of the QCCNN across different input scaling factors is only conducted on our bigger dataset – OrganAMNIST – utilizing one encoding method – the higher-order encoding – due to its superior performance in prior research [15].

Fig. 2 shows the evolution of the best training and best validation accuracies achieved within 20 training epochs for scaling factors in the encoder ranging from  $f = \pi/4$  to  $f = 4\pi$ . Note that at  $\pi$  and given the normalization, the rotation gates already cover the full circle. However, for experimental purposes, we extend to  $4\pi$ . This figure clearly shows that the validation accuracy drops dramatically as the scaling factor grows. While this could be expected for values above  $\pi$ , this result is surprising, as one would expect the performance to be best when the coverage of the Hilbert space of the circuit is high. Further interpretation on this phenomenon will be provided in section IV-C. Interestingly, a good training accuracy can nevertheless be obtained with all tested scaling factors. In the rest of this paper, a scaling factor of  $\pi/4$  was selected for all encodings to maximize the obtained accuracies.




Fig. 2: Comparison of the performance of a hybrid QCCNN for different scaling factors on OrganAMNIST using the higher-order encoding.

2) *Performance with varying encodings:* We proceed to compare the performance of the developed QCCNNs for the chosen encoding methods using data-reuploading. The performances are analyzed in terms of best reached training and validation accuracies achieved within 20 training epochs. For every QCCNN architecture, three training sessions were conducted with different weight initialization. The reported metrics represent the averages across these training runs, and the error bands reflect the variance observed. Fig. 3 compares the accuracies achieved by the models with the  $R_X$ ,  $R_Y$ , higher-order, and amplitude encodings for 1 up to 5 reuploading layers on BreastMNIST and OrganAMNIST.

With one encoding layer, the higher-order and  $R_X$  encoders achieve the best validation accuracy on both datasets, even when taking into account the error bands that are notably larger on the BreastMNIST dataset, due to the smaller size of the dataset. The  $R_Y$  encoding performs a little bit worse, while the amplitude encoding performs significantly worse in validation accuracy, despite achieving a training accuracy on-par with the other models. This gap in performance of the architecture with amplitude encoding is likely attributed to its utilization of a smaller number of qubits (2 instead of 4 as seen in other approaches).

With two layers, the higher-order encoder exhibits a slight decline compared to other methods, while the  $R_X$  encoder remains stagnant, and the  $R_Y$  encoder achieves validation accuracies comparable to (on BreastMNIST) or even surpasses the two previous methods (on OrganAMNIST). The amplitude encoder benefits from the additional reuploading layer but the performance remains less than ideal.

When adding three and more reuploading layers, the gain and loss in accuracy due to the additional layers seems to saturate as the accuracies remain more or less constant on Or-

ganAMNIST for all architectures but the amplitude encoding. In particular, they reach a similar final accuracy around 87%. On BreastMNIST, in spite of the relatively large error bands, we observe a similar behavior of saturation of the performance. On the other hand, the architecture with amplitude encoding alternatively varies as the number of reuploading layers increases, while overall improving but remaining at a low value compared to the other encoders.

All in all, the highest mean validation accuracy in all considered models is achieved on both datasets with the  $R_Y$  encoding, at 4 and 2 layers for the BreastMNIST and OrganAMNIST, respectively.

### C. Quantum metrics and their correlation to performance

We now examine the various metrics outlined in section III-C and explore potential correlations with the performance presented in IV-B2. The highest validation accuracy attained for all models is plotted against the analyzed quantum metrics in Fig. 4. We first note that these correlation plots look similar for both datasets.

Starting with the expressibility, we see that most circuits exhibit low values, making them highly expressible, with the exception of the  $R_X$ ,  $R_Y$  and higher-order encodings with one layer, which have comparatively higher values. Increasing the number of data-reuploading layers results in an increase in the expressibility, but is not necessarily accompanied with an improved validation accuracy. Among the more expressible circuits, we encounter instances with good accuracy such as the four-layer  $R_Y$  encoding, while others, such as amplitude encoding with 4 layers, demonstrate comparatively poorer performance. This tends to indicate that there is no clear correlation between the best achieved accuracy and the expressibility.

It should additionally be noted that in the experiment outlined in section IV-B1, where a fixed encoding strategy, the higher-order encoding, was employed, the accuracy notably decreased with an increase in the scaling factor. However, a higher scaling factor corresponds to a more expressible circuit as it allows it to encompass a larger portion of the Hilbert space. This observation suggests a negative correlation between expressibility and performance achieved. Such an effect may stem from training challenges associated with high expressibility, as discussed in [30]. The study highlights that low-expressibility ansätze (characterized by a high expressibility value) can result in both small and large cost gradients, whereas high-expressibility ansätze (low expressibility value) tend to exhibit predominantly flat cost landscapes, making them generally challenging to train.

Next, we examine the entanglement capacity of the various encodings. One first notes that the amplitude encoding displays clearly lower entangling capabilities than other setups. This is expected as only 2 qubits are used in this case. For the other architectures, one observes a cluster around high values of entanglement capacity, with the exception of the higher-order encoding with one layer. No correlation is therefore observable in the plot between the best validation accuracy

and the entanglement capability. For instance, both the four-layer  $R_Y$  and the two-layer higher-order encoding settings exhibit similar metric values but show significantly divergent performances across both datasets.


Lastly, the normalized effective dimension appears to follow a downward trend as the number of layers increases, with the exception of the  $R_Y$  encoding, where an increase in NED is observed from layer 1 to layer 2, coinciding with a significant increase in validation accuracy. Although no clear correlation is observed between the NED and the VQCs' performances, these results seem to indicate a saturation effect as the number of layers increases, given four qubits and the chosen application datasets, as suggested in [21].

### D. Fourier Coefficients and their correlation to performance


In the following subsection, the Fourier coefficients generated by the angle encodings and the higher-order encoding are examined, and potential relations to the performance are discussed. The amplitude encoding is not addressed here due to the state normalization requirement. To address visualization challenges associated with higher-dimensional cases, this analysis is done in a univariate scenario, where the same input is passed to each qubit in the quantum circuit. The output is classically decomposed into its constituent frequencies using the discrete Fourier decomposition. The coefficients corresponding to the various frequencies can then be displayed on the complex plane-where the x-axis represents the real part and the y-axis is the imaginary component-and analyzed. This study is a simplification of what really happens in the quantum convolution, where at each iteration of the convolution process, one pixel value is mapped onto one qubit, and where therefore all mapped values differ. This univariate analysis neglects the possibility of having different input pixel values at a given stride of the filter. Nonetheless, it provides some insights into the expressive power of the different circuits.

1) *Experimental procedure:* A set of 21 ordered, linearly spaced inputs within  $[-1, 1]$  are sampled and passed onto the circuit. This replicates the possible input values in the normalized datasets. 100 weights are additionally randomly sampled from  $[0, 2\pi[$  from a uniform distribution. The outputs are then measured and decomposed using the Real Fast Fourier Transform (RFFT) implemented in the Numpy package [31]. The chosen number of coefficients to be considered was set to 10, plus the zero frequency coefficient. The distribution of the resulting coefficients is displayed in the complex plane. For each encoding in this analysis, the procedure is performed for up to four data reuploading layers. The effect of the chosen embedding strategy as well as the increase in the number of layers is studied.

Following the analysis presented in [21] and by putting equations 12 and 13 together, the number of adaptable degrees of freedom is expected to saturate at two circuit layers for the two angle encodings (where the dimensionality of the local system  $n = 1$ ), and at four layers for the higher-order encoding (where  $n = 2$  due to the pairwise CNOT operations in the encoding).



(a) On the BreastMNIST dataset



(b) On the OrganAMNIST dataset

Fig. 3: Comparison of the performance of 4 encoding methods for up to 5 layers for hybrid QCCNNs in terms of best training and validation accuracies.

2) *Fourier coefficients analysis*: The distributions of the Fourier coefficients for each qubit for the  $R_X$ ,  $R_Y$  and higher-order encodings are shown in Fig. 5. We particularly pay attention to the number of non-null frequencies present, their rank (from 0 to 10) and the distribution of their coefficients.

At one layer, we first observe that more frequencies are present in the  $R_X$  and higher-order encodings compared to the  $R_Y$  encoding, as for  $R_Y$ , the number of non-null Fourier coefficients is limited to the first 4, as opposed to all 10 for the  $R_X$  encoding and also 10 for the higher-order encoding on qubits 1 and 3. In terms of their coefficient, while a mix of both real and imaginary is accessible with the  $R_X$  and higher-order encodings (both on the x and y-axis), the  $R_Y$  encoding is limited to purely real values (only on the x-axis). When comparing to the performance of the QCCNNs, we see that on both datasets,  $R_Y$  performs worse than  $R_X$  and the higher-order encodings at 1 layer. This seems to be in line with the interpretation of the expressibility of the Fourier coefficients.


In the  $R_X$  angle encoding, the addition of one data-reuploading layer results in a more widespread distribution of the first coefficient at qubit 2. As the number of layers is further increased, the number of accessible frequencies is reduced in qubit 4 and then in qubits 2 and 3. These findings are consistent with [21], where it is suggested that the number of adjustable degrees of freedom will saturate at two circuit layers. In comparison, the best accuracy is achieved at 2 layers for the BreastMNIST dataset, followed with a drop in accuracy at 3 layers, which aligns again with these results. On the OrganAMNIST dataset, the best performance is attained at 3 layers but with only a marginal improvement observed from the second to the third layer.

For  $R_Y$  and the higher-order encoding, increasing the number of data-reuploading layers causes the variance of the


Fourier coefficients to increase and their values to become more mixed (real and imaginary). For both encodings, this increase is mostly noticeable when transitioning from one to two circuit layers, and becomes less visible when increasing further. This is an indication of a saturation of adaptable degrees of freedom at two layers, which was expected for  $R_Y$ , but was expected at 4 layers for the higher-order encoding. Regarding performance, the best performance with  $R_Y$  occurs at 4 and 2 layers on BreastMNIST and OrganAMNIST, respectively. The findings on OrganAMNIST are in line with both the observed distribution of Fourier coefficients and the previous work on the saturation of adjustable degrees of freedom. However, the result on BreastMNIST diverges, likely due to the limited size of the dataset. For the higher-order encoding, the best outcomes are achieved at 3 and 4 layers, indicating the relevance of the analysis in [21].

## V. CONCLUSION


We analyze the performance of various hybrid quantum-classical convolutional neural networks composed of different encodings on two medical imaging tasks, namely breast cancer detection on ultrasound images and organ classification with CT scans. We first observe that the encoding used – be it  $R_X$ ,  $R_Y$  angle encoding, higher-order, or amplitude encoding – leads to significantly different performances, highlighting the critical significance of encoding selection when implementing quantum variants of classical algorithms. We additionally examine the effects of input scaling and data-reuploading and show that a well-chosen parametrization of these two elements can also lead to a substantial gain in performance on a chosen dataset. These findings point to the need for continued investigation into QCCNNs, considering their compatibility with current NISQ hardware thanks to the minimal qubit demands and




(a) Expressibility on the BreastMNIST dataset




(b) Expressibility on the OrganAMNIST dataset




(c) Entangling capability on the BreastMNIST dataset



(d) Entangling capability on the OrganAMNIST dataset




(e) NED on the BreastMNIST dataset




(f) NED on the OrganAMNIST dataset


Fig. 4: Correlation plots between the performance of 4 encoding methods for up to 5 layers and the expressibility, the entangling capability and the normalized effective dimension. The numerical labels indicate the number of layers.




(1.a) 1 layer



(1.b) 2 layers




(1.c) 3 layers




(1.d) 4 layers


(1)  $R_X$  encoding




(2.a) 1 layer



(2.b) 2 layers




(2.c) 3 layers




(2.d) 4 layers


(2)  $R_Y$  encoding




(3.a) 1 layer



(3.b) 2 layers



(3.c) 3 layers



(3.d) 4 layers

(3) Higher-order encoding

Fig. 5: Distributions of the Fourier coefficients of the outputs generated by the various encodings and basic entangling circuit for different reuploading layers.

shallow quantum circuit depths involved. When attempting to explain these performances using quantum metrics such as the expressibility, the entanglement capability, and the normalized effective dimension of the circuit, a clear correlation is yet to be observed. On the other hand, a discernible trend begins to emerge upon closer examination of the Fourier coefficients generated by the quantum circuits, where a high number of non-null frequency as well as a wide-spread distribution of coefficients across both real and imaginary axes is linked to a good performance. Moreover, one can observe changes in the distribution of coefficients, thus indicating potential saturation effects as additional reuploading layers are incorporated. Still, additional work is needed to understand how these findings can effectively inform encoding selection based on the dataset employed. Furthermore, future work should also focus on investigating the impact of noise on these architecture, via simulation packages, but also by running the results on the hardware to comprehensively assess its impact.

## VI. ACKNOWLEDGEMENTS

The project/research is supported by the Bavarian Ministry of Economic Affairs, Regional Development and Energy with funds from the Hightech Agenda Bayern.

## REFERENCES

- [1] M. C. Caro, H.-Y. Huang *et al.*, “Generalization in quantum machine learning from few training data,” *Nature Communications*, vol. 13, no. 1, p. 4919, Aug. 2022, arXiv:2111.05292 [quant-ph, stat]. [Online]. Available: <http://arxiv.org/abs/2111.05292>
- [2] I. Cong, S. Choi, and M. D. Lukin, “Quantum Convolutional Neural Networks,” *Nature Physics*, vol. 15, no. 12, pp. 1273–1278, Dec. 2019, arXiv:1810.03787 [cond-mat, physics:quant-ph]. [Online]. Available: <http://arxiv.org/abs/1810.03787>
- [3] M. Schuld and N. Killoran, “Quantum machine learning in feature Hilbert spaces,” *Physical Review Letters*, vol. 122, no. 4, p. 040504, Feb. 2019, arXiv:1803.07128 [quant-ph]. [Online]. Available: <http://arxiv.org/abs/1803.07128>
- [4] V. Havlicek, A. D. Córcoles *et al.*, “Supervised learning with quantum enhanced feature spaces,” *Nature*, vol. 567, no. 7747, pp. 209–212, Mar. 2019, arXiv:1804.11326 [quant-ph, stat]. [Online]. Available: <http://arxiv.org/abs/1804.11326>
- [5] D. Mattern, D. Martyniuk *et al.*, “Variational Quantum Convolutional Neural Networks with enhanced image encoding,” Jun. 2021, arXiv:2106.07327 [cs]. [Online]. Available: <http://arxiv.org/abs/2106.07327>
- [6] A. Pérez-Salinas, A. Cervera-Lierta *et al.*, “Data re-uploading for a universal quantum classifier,” *Quantum*, vol. 4, p. 226, Feb. 2020, arXiv:1907.02085 [quant-ph]. [Online]. Available: <http://arxiv.org/abs/1907.02085>
- [7] M. Schuld, R. Sweke, and J. J. Meyer, “The effect of data encoding on the expressive power of variational quantum machine learning models,” *Physical Review A*, vol. 103, no. 3, p. 032430, Mar. 2021, arXiv:2008.08605 [quant-ph, stat]. [Online]. Available: <http://arxiv.org/abs/2008.08605>
- [8] I. Kerenidis, J. Landman, and A. Prakash, “Quantum Algorithms for Deep Convolutional Neural Networks,” Nov. 2019, arXiv:1911.01117 [quant-ph]. [Online]. Available: <http://arxiv.org/abs/1911.01117>
- [9] T. Hur, L. Kim, and D. K. Park, “Quantum convolutional neural network for classical data classification,” *Quantum Machine Intelligence*, vol. 4, no. 1, p. 3, Jun. 2022, arXiv:2108.00661 [quant-ph]. [Online]. Available: <http://arxiv.org/abs/2108.00661>
- [10] Y. Li, R.-G. Zhou *et al.*, “A quantum deep convolutional neural network for image recognition,” *Quantum Science and Technology*, vol. 5, no. 4, p. 044003, Jul. 2020. [Online]. Available: <https://iopscience.iop.org/article/10.1088/2058-9565/ab9f93>
- [11] S. Wei, Y. Chen *et al.*, “A Quantum Convolutional Neural Network on NISQ Devices,” Apr. 2021, arXiv:2104.06918 [quant-ph]. [Online]. Available: <http://arxiv.org/abs/2104.06918>
- [12] Y. Lü, Q. Gao *et al.*, “A Quantum Convolutional Neural Network for Image Classification,” in *2021 40th Chinese Control Conference (CCC)*, Jul. 2021, pp. 6329–6334, iSSN: 1934-1768. [Online]. Available: <https://ieeexplore.ieee.org/document/9550027>
- [13] M. Henderson, S. Shakyia *et al.*, “Quantum Convolutional Neural Networks: Powering Image Recognition with Quantum Circuits,” Apr. 2019, arXiv:1904.04767 [quant-ph]. [Online]. Available: <http://arxiv.org/abs/1904.04767>
- [14] J. Liu, K. H. Lim *et al.*, “Hybrid Quantum-Classical Convolutional Neural Networks,” *Science China Physics, Mechanics & Astronomy*, vol. 64, no. 9, p. 290311, Sep. 2021, arXiv:1911.02998 [quant-ph]. [Online]. Available: <http://arxiv.org/abs/1911.02998>
- [15] A. Matic, M. Monnet *et al.*, “Quantum-classical convolutional neural networks in radiological image classification,” Aug. 2022, arXiv:2204.12390 [quant-ph]. [Online]. Available: <http://arxiv.org/abs/2204.12390>
- [16] M. Monnet, H. Gebran *et al.*, “Pooling techniques in hybrid quantum-classical convolutional neural networks,” May 2023, arXiv:2305.05603 [quant-ph]. [Online]. Available: <http://arxiv.org/abs/2305.05603>
- [17] S. Sim, P. D. Johnson, and A. Aspuru-Guzik, “Expressibility and entangling capability of parameterized quantum circuits for hybrid quantum-classical algorithms,” *Advanced Quantum Technologies*, vol. 2, no. 12, p. 1900070, Dec. 2019, arXiv:1905.10876 [quant-ph]. [Online]. Available: <http://arxiv.org/abs/1905.10876>
- [18] A. Abbas, D. Sutter *et al.*, “The power of quantum neural networks,” *Nature Computational Science*, vol. 1, no. 6, pp. 403–409, Jun. 2021, arXiv:2011.00027 [quant-ph]. [Online]. Available: <http://arxiv.org/abs/2011.00027>
- [19] T.-A. Drăgan, M. Monnet *et al.*, “Quantum Reinforcement Learning for Solving a Stochastic Frozen Lake Environment and the Impact of Quantum Architecture Choices,” Dec. 2022, arXiv:2212.07932 [quant-ph]. [Online]. Available: <http://arxiv.org/abs/2212.07932>
- [20] J. G. Vidal and D. O. Theis, “Input Redundancy for Parameterized Quantum Circuits,” Jun. 2020, arXiv:1901.11434 [quant-ph]. [Online]. Available: <http://arxiv.org/abs/1901.11434>
- [21] B. Casas and A. Cervera-Lierta, “Multidimensional Fourier series with quantum circuits,” *Physical Review A*, vol. 107, no. 6, p. 062612, Jun. 2023, arXiv:2302.03389 [quant-ph]. [Online]. Available: <http://arxiv.org/abs/2302.03389>
- [22] Aikaterini, Gratsea, and P. Huembeli, “The effect of the processing and measurement operators on the expressive power of quantum models,” Nov. 2022, arXiv:2211.03101 [quant-ph]. [Online]. Available: <http://arxiv.org/abs/2211.03101>
- [23] M. Mottonen, J. J. Vartiainen *et al.*, “Transformation of quantum states using uniformly controlled rotations,” Jul. 2004, arXiv:quant-ph/0407010. [Online]. Available: <http://arxiv.org/abs/quant-ph/0407010>
- [24] V. Bergholm, J. Izaac *et al.*, “PennyLane: Automatic differentiation of hybrid quantum-classical computations,” Jul. 2022, arXiv:1811.04968 [physics, physics:quant-ph]. [Online]. Available: <http://arxiv.org/abs/1811.04968>
- [25] K. Zyczkowski and H.-J. Sommers, “Average fidelity between random quantum states,” *Physical Review A*, vol. 71, no. 3, p. 032313, Mar. 2005, arXiv:quant-ph/0311117. [Online]. Available: <http://arxiv.org/abs/quant-ph/0311117>
- [26] A. Paszke, S. Gross *et al.*, “PyTorch: An Imperative Style, High-Performance Deep Learning Library,” Dec. 2019, arXiv:1912.01703 [cs, stat]. [Online]. Available: <http://arxiv.org/abs/1912.01703>
- [27] J. Yang, R. Shi, and B. Ni, “MedMNIST Classification Decathlon: A Lightweight AutoML Benchmark for Medical Image Analysis,” in *2021 IEEE 18th International Symposium on Biomedical Imaging (ISBI)*, Apr. 2021, pp. 191–195, arXiv:2010.14925 [cs]. [Online]. Available: <http://arxiv.org/abs/2010.14925>
- [28] J. Yang, R. Shi *et al.*, “MedMNIST v2 – A large-scale lightweight benchmark for 2D and 3D biomedical image classification,” *Scientific Data*, vol. 10, no. 1, p. 41, Jan. 2023, arXiv:2110.14795 [cs, eess]. [Online]. Available: <http://arxiv.org/abs/2110.14795>
- [29] P. Bilic, P. Christ *et al.*, “The Liver Tumor Segmentation Benchmark (LiTS),” *Medical Image Analysis*, vol. 84, p. 102680, Feb. 2023, arXiv:1901.04056 [cs]. [Online]. Available: <http://arxiv.org/abs/1901.04056>
- [30] Z. Holmes, K. Sharma *et al.*, “Connecting ansatz expressibility to gradient magnitudes and barren plateaus,” *PRX Quantum*, vol. 3, no. 1, p. 010313, Jan. 2022, arXiv:2101.02138 [quant-ph, stat]. [Online]. Available: <http://arxiv.org/abs/2101.02138>

- [31] C. R. Harris, K. J. Millman *et al.*, “Array Programming with NumPy,” *Nature*, vol. 585, no. 7825, pp. 357–362, Sep. 2020, arXiv:2006.10256 [cs, stat]. [Online]. Available: <http://arxiv.org/abs/2006.10256>