



Building Continuous Quantum-Classical Bayesian Neural Networks for a Classical Clinical Dataset

Alona Sakhnenko
Fraunhofer IKS
Munich, Bavaria, Germany
alona.sakhnenko@iks.fraunhofer.de

Julian Sikora
Fraunhofer IKS
Technical University of Munich
Munich, Bavaria, Germany
julian.sikora@iks.fraunhofer.de

Jeanette Miriam Lorenz
Fraunhofer IKS
Ludwig-Maximilian University
Munich, Bavaria, Germany
jeanette.miriam.lorenz@iks.fraunhofer.de

ABSTRACT

In this work, we are introducing a Quantum-Classical Bayesian Neural Network (QCBNN) that is capable to perform uncertainty-aware classification of classical medical dataset. This model is a symbiosis of a classical Convolutional NN that performs ultra-sound image processing and a quantum circuit that generates its stochastic weights, within a Bayesian learning framework. To test the utility of this idea for the possible future deployment in the medical sector we track multiple behavioral metrics that capture both predictive performance as well as model's uncertainty. It is our ambition to create a hybrid model that is capable to classify samples in a more uncertainty aware fashion, which will advance the trustworthiness of these models and thus bring us step closer to utilizing them in the industry. We test multiple setups for quantum circuit for this task, and our best architectures display bigger uncertainty gap between correctly and incorrectly identified samples than its classical benchmark at an expense of a slight drop in predictive performance. The innovation of this paper is two-fold: (1) combining of different approaches that allow the stochastic weights from the quantum circuit to be continues thus allowing the model to classify application-driven dataset; (2) studying architectural features of quantum circuit that make-or-break these models, which pave the way into further investigation of more informed architectural designs.

CCS CONCEPTS

• **Computing methodologies** → **Supervised learning; Neural networks**; • **Hardware** → **Quantum computation**; • **Mathematics of computing** → **Bayesian computation**.

KEYWORDS

Quantum Bayesian Neural Networks, QCBNN, Hybrid models, PQC architecture study

ACM Reference Format:

Alona Sakhnenko, Julian Sikora, and Jeanette Miriam Lorenz. 2024. Building Continuous Quantum-Classical Bayesian Neural Networks for a Classical Clinical Dataset. In *Recent Advances in Quantum Computing and Technology*



This work is licensed under a Creative Commons Attribution International 4.0 License.

ReAQCT '24, June 19–20, 2024, Budapest, Hungary
© 2024 Copyright held by the owner/author(s).
ACM ISBN 979-8-4007-1798-7/24/06
<https://doi.org/10.1145/3665870.3665872>

(ReAQCT '24), June 19–20, 2024, Budapest, Hungary. ACM, New York, NY, USA, 11 pages. <https://doi.org/10.1145/3665870.3665872>

1 INTRODUCTION

Neural Networks (NNs) have achieved incredible results in many fields of Machine Learning. However, some practical applications, such as medical AI, have strict reliability requirements where NNs may lead to catastrophic failure if they cannot reliably estimate the confidence interval of their output. Bayesian learning provides a well-established framework to train and analyse uncertainty-aware models, and Bayesian NNs (BNNs) are stochastic NNs that yield the benefits of Bayesian learning in a deep learning setup. The success of these models depend on a good choice of both a NN architecture and stochastic distribution family.

Quantum Computing (QC) is an emerging technology that has been rapidly gaining momentum in the last years. Many research initiatives have been directed towards investigating early applications of this technology in a real-world scenario. There is, however, one substantial challenge for application-driven approaches, which is Noisy Intermediate-Scale Quantum (NISQ) devices that are available today. Unlike fault-tolerant computers that promise provable advantages, such as with Shor and HHL algorithms, the advantages of NISQ algorithms are still being explored.

Early proof-of-concept applications of NISQ for Machine Learning (ML) problems have shown a lot of potential of boosting capacities of ML models. These applications have provided both *theory-driven* insights, e.g. Quantum Neural Networks showed higher effective dimensions on an investigated problem [1], Quantum Kernels allow for better generalization on specific datasets [7], Born Machines are capable of representing distributions that are hard to sample classically [5], as well as *empirically-driven* results, e.g. boosted performance of a hybrid Quantum-Classical Generative NN for high-resolution images [16], boosted performance of a hybrid Quantum-Classical Convolutional NN (QCCNN) on medical dataset while requiring less training parameters than the classical benchmark [10].

In this work, we explore how we can leverage NISQ devices to create uncertainty-aware models by utilizing quantum circuits for Bayesian learning and how this setup compares to its classical analog. We concentrate on application-driven approach that is suitable for potential future deployment in medical AI sector. As a proof-of-concept use-case we select a publicly available Breast Ultrasound Images (BreastMNIST) dataset. We combine several architectural ideas together, e.g. *binary* quantum-classical BNN [13], *continuous* quantum Generative Adversarial Networks [15] that propose

an idea to sample beyond the discrete distribution as well as QC-CNN [10] that proved useful in a hybrid quantum-classical setting on the same dataset. We create a Quantum-Classical (Convolutional) BNN (QCBNN) that processes ultrasound images classically, while the stochastic continuous weights are generated by a quantum circuit. We track multiple behavioral metrics of the model that capture both its predictive performance and its uncertainty. The proposed model shows a bigger gap in confidence of correctly and incorrectly identified samples than its classical benchmark, which means our model showed higher uncertainty on datapoints that were misclassified and hence increasing the trustworthiness of these models. Furthermore, we perform a systematic study of the effects of different parametrized quantum circuit (PQC) architectures from the literature, and build tailor-made PQCs for this task that increase the confidence gap. We provide the analysis of these architectures from a standpoint of different metrics, including confidence errors and ensemble sizes of correctly and incorrectly identified samples, model’s calibration and the shape of weights distribution, which offers insights on how different architectural elements influence models’ behaviour.

2 RELATED WORK

One of the early explorations into utility of fault-tolerant QC for Bayesian learning has been performed by *Zhao et al.* [20]. This work proposed to use HHL algorithm to accelerate matrix inversion procedure in the context of Gaussian Processes (GP). GPs have many beneficial properties, but, due to scaling bottlenecks, these models are not often used in practice. HHL algorithm can be leveraged to speed up the inference procedure of these models and, hence, improve their applicability. *Berner et al.* [3] proposed to utilize another fault-tolerant algorithms for inner product estimation to accelerate training and inference of BNNs. Both of these algorithms have shown their usability on small examples, however, making them suitable for real-world instances would require either viable NISQ-adaptations of these algorithms or fault-tolerant QCs.

Other works looked into Bayesian learning in NISQ era. Work by *Nguyen and Chen* [12] showed the utility of Bayesian learning for training Quantum NNs. The authors showed that their method yield a better generalization (based on the study of model’s capacity), and provided additional information on epistemic uncertainty. In our work, we investigate Bayesian learning for stochastic models in contrast to deterministic models [12], which is a more common method for building uncertainty-aware models. One of the first attempts to translate stochastic models into the quantum realm was done by *Nikoloska and Simeone* [13], where the authors introduced a quantum binary BNN that utilizes a Born machine to generate binary stochastic weights. The authors show that this approach allows for an easier way to optimize over discrete distributions, and the utility of it was shown on a simple regression example. Scaling this idea to more complex problems might prove difficult because of limited capacity of binary NNs. Hence, in this work we investigate how we can extend this approach to sample from continuous distributions and its potential benefits on a medical use-case.

Matic et al. [10] demonstrated a benefit of using quantum-classical setup to boost performance on the same BreastMNIST dataset. In

their work, the authors propose a Quantum-Classical Convolutional NNs, in which PQC performed a convolution operation and the expectation values of each qubit were propagated further into a fully-connected layer. This model achieved a better performance than its classical benchmark while requiring less trainable parameters. Subsequent work [11] tested additionally different quantum pooling techniques, which allowed to boost performance even further. In our work, we build upon the architecture that was used in [10], but change the role of a PQC to sample from weights distribution instead of performing convolution operation. We concentrate on the most successful PQC architectures to test their utility for our task. In our work, we extend the family of investigated PQCs and investigate which of their features are beneficial in our setup.

3 BACKGROUND

In this section, we provide a detailed description of important notions from classical (Section 3.1) and quantum machine learning (Section 3.2-3.3) that are essential building blocks of QCBNN.

3.1 Bayesian Neural Networks

Bayesian Neural Networks (BNNs) are stochastic NN trained using Bayesian inference, which comprise of (1) a stochastic model of a chosen parametrization prior $p(\mathbf{w})$ and a prior of confidence $p(y|x, \mathbf{w})$, which result in a following posterior on training data D :

$$p(\mathbf{w}|D) = \frac{p(D_y|D_x, \mathbf{w})p(\mathbf{w})}{\int_{\mathbf{w}'} p(D_y|D_x, \mathbf{w}')p(\mathbf{w}')d\mathbf{w}'}, \quad (1)$$

and (2) a functional model $\Phi_{\mathbf{w}}(x)$, which in this case is a NN. Prediction process with BNNs can be represented in a general case as [9]:

$$\begin{aligned} \mathbf{w} &\sim p(\mathbf{w}|D) \\ \mathbf{y} &= \Phi_{\mathbf{w}}(\mathbf{x}) + \epsilon, \end{aligned} \quad (2)$$

where \mathbf{x} is an input, \mathbf{y} is a corresponding label and ϵ is an approximation error between output for a functional model output and the label. To perform classification with BNNs, we can compute an average model prediction over the ensemble of N different outputs. To begin with, we compute the probability of each class as follows:

$$\hat{\mathbf{p}} = \frac{1}{N} \sum_{i=0}^N \Phi_{\mathbf{w}_i}(\mathbf{x}), \quad (3)$$

from which we can compute the final prediction as $\hat{\mathbf{y}} = \arg \max_i p_i \in \hat{\mathbf{p}}$.

The challenge in working with BNNs resides in the fact that sampling directly from the posterior distribution (Equation 1) is difficult due to intractability of the evidence $\int_{\mathbf{w}'} p(D_y|D_x, \mathbf{w}')p(\mathbf{w}')d\mathbf{w}'$ [9]. Instead, most commonly, we can either build a Markov chain with a stationary distribution being the posterior (Markov chain Monte Carlo methods) or select a tractable distribution and adjust its parameters to approximate the posterior (Variational inference). Due to its scalability, Variational inference (VI) has gain popularity within the ML community and is therefore the focus of this work. For VI we select a family of tractable distributions $q_{\theta}(\mathbf{x}) \in \mathcal{Q}$ and tweak the variational parameters θ to minimize the Kullback–Leibler (KL) divergence between two distributions:

$$\mathbb{KL}[q_\theta(\mathbf{w}|D)||p(\mathbf{w}|D)] = \mathbb{E}_{\mathbf{w} \sim q_\theta(\mathbf{w}|D)} \left[\log \frac{q_\theta(\mathbf{w}|D)}{p(\mathbf{w}|D)} \right] \quad (4)$$

3.2 Variational Inference and Adversarial Training on a QC

Born machines offer a way to model a distribution using a quantum state $|\psi(\theta, x)\rangle$, that encodes an input vector x and is parameterized by θ . These models generate bitstrings with probabilities $q_\theta(\mathbf{w}|x) = |\langle \mathbf{w} | \psi(\theta, x) \rangle|^2$ [2]. The power of this method to model intractable distributions has been shown in [4], and its generalization capabilities have been investigated in [6].

A VI training loop for Born machines was proposed by *Benedetti et al.* [2]. As mentioned above, KL divergence is a popular objective function for VI, which we can simplify as follows:

$$\begin{aligned} \mathbb{KL}[q_\theta(\mathbf{w}|D)||p(\mathbf{w}|D)] & \quad (5) \\ &= \mathbb{E}_{\mathbf{w} \sim q_\theta(\mathbf{w}|D)} \left[\log \frac{q_\theta(\mathbf{w}|D)}{p(\mathbf{w}|D)} \right] \\ &= \mathbb{E}_{\mathbf{w} \sim q_\theta(\mathbf{w}|D)} \left[\log \frac{q_\theta(\mathbf{w}|D)p(D)}{p(D|\mathbf{w})p(\mathbf{w})} \right] \\ &= \mathbb{E}_{\mathbf{w} \sim q_\theta(\mathbf{w}|D)} \left[\log \frac{q_\theta(\mathbf{w}|D)}{p(\mathbf{w})} - \log p(D|\mathbf{w}) \right] + \text{const.} \quad (6) \end{aligned}$$

Unfortunately, we cannot estimate $\frac{q_\theta(\mathbf{w}|D)}{p(\mathbf{w})}$ directly on a quantum computer. To tackle this issue, adversarial learning can be utilized. We can construct a binary classifier $d_\phi(\mathbf{w}, D)$, which discriminates between samples from the true distribution and Born machine generated ones. The classifier objective function is defined as a cross-entropy:

$$\mathcal{G}_{KL}(\phi; \theta) = \mathbb{E}_{\mathbf{w} \sim q_\theta(\mathbf{w}|D)} [\log d_\phi] + \mathbb{E}_{\mathbf{w} \sim q_\theta(\mathbf{w})} [\log(1 - d_\phi)] \quad (7)$$

In order to extract the odds of a data point coming from one of the two distributions, the classifier's output can be transformed using the logit function:

$$\text{logit}(d_\phi(\mathbf{w}, D)) = \log \frac{d_\phi(\mathbf{w}, D)}{1 - d_\phi(\mathbf{w}, D)} \approx \log \frac{q_\theta(\mathbf{w}, D)}{p(\mathbf{w})}. \quad (8)$$

Therefore, the final objective of the Born training is:

$$\mathcal{L}_{KL}(\theta; \phi) = \mathbb{E}_{\mathbf{w} \sim q_\theta(\mathbf{w}|D)} [\text{logit}(d_\phi(\mathbf{w}, D)) - \log p(D|\mathbf{w})]. \quad (9)$$

During the overall training procedure, both a classifier and a Born machine are trained simultaneously as follows:

$$\begin{aligned} & \max_{\phi} \mathcal{G}_{KL}(\phi; \theta) \\ & \min_{\theta} \mathcal{L}_{KL}(\theta; \phi). \quad (10) \end{aligned}$$

This adversarial training strategy is akin to the one commonly used for Generative Adversarial Networks (GANs). Quantum adaptation of these networks have been successfully applied to a variety of different use-cases [8, 16, 21]. *Romero and Aspuru-Guzik* [15] presented an approach that allows a quantum GANs to sample from a continuous distribution. Their approach suggests to inject classical noise into a PQC and a set of observables, whose expectation values serve as an output of the model. This method can be used to create continuous distributions for quantum VI method described above.

3.3 Quantum-Classical Bayesian Neural Networks

As mentioned in Section 2, *Nikoloska and Simeone* [13] proposed an initial architecture for QCBNNs. Their method followed the same techniques described in Section 3.2 and generated a binary distribution that proved its viability on a toy regression example in meta learning setting. The loss function of the model comprise of likelihood and adversarial terms as follows:

$$\mathcal{L} = -\alpha \mathbb{E}[\log p(D|\theta)] + \beta \mathcal{L}_{KL}, \quad (11)$$

where $\alpha, \beta > 0$ are regularizing hyperparameters.

4 METHOD

In this project, we build a QCBNN that is capable to process a real-world medical datasets, such as ultrasound images known as BreastMNIST [19], as well as deliver reliable uncertainty estimations. In the following text, we outline the details of our implementation and specific performance metrics that are used to evaluate the models.

4.1 Use-case

We focus on BreastMNIST [19] dataset, which consists of breast ultrasound 28×28 images that are classified into malignant and non-malignant categories. This dataset is interesting because it captures the complexity of identifying cancerous samples in a scarce unbalanced setting, as it consists of a mere 780 samples for training, validation and testing. The data is normalized as the only further preprocessing step.

4.2 Architecture

We start with an architectural idea that proved useful in a hybrid quantum-classical setup on this dataset, namely QCCNN [10] (see Section 2). We create a Quantum-Classical (Convolutional) BNN (QCBNN) by changing the function of a PQC from performing a convolutional operation to sampling stochastic weights for the classically performed convolution. A sketch of a full architecture is presented in Fig. 1. The training strategy is adapted from binary quantum BNN as described in Section 3.3. The model consists of a quantum *generator* for sampling weights and trainer modules, such as a *prior* and a *discriminator*, that attempts to distinguish quantumly generated samples from the prior. To boost expressivity of the model we adapt a technique for sampling from continuous weights described in Section 3.2, which introduces a classical *noise* and a *postprocessing* layer to our architecture.

The PQC structure in this setup requires two architectural decisions: an *embedding layer* that encodes a classical noise vector and a *calculation layer*. Here, we select an embedding strategy that showed the most promise in other works [1, 10] and our internal experiments, namely a *higher-order embedding* layer. The structures of calculation layer, however, are still heavily debated in the literature. Hence, we test out multiple variations of these layers to see which architectural features of PQCs have the most impact on the performance of the overall model. We first test models that were presented in the literature, such as *Matic I* (Fig. 2a) and *Matic II* (Fig. 2b) (the best and the second best architectures from [10]), *Nikoloska* (Fig. 2c) inspired by the PQC from binary QBNN

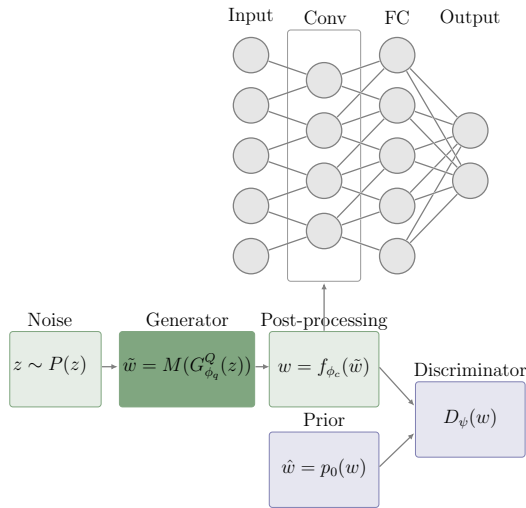


Figure 1: General architecture of proposed continuous hybrid quantum-classical BNN model. The top gray part represents a classical CNN architecture, while the bottom green part represent the quantum weights sampler and the bottom blue part represents the trainers for the quantum sampler. The quantum sampler comprises multiple modules: classical noise, a PQC generator and a post-processing step that computes expectation values of each qubit. The only part that requires a quantum hardware is the quantum generator, while all the other components are performed classically. The trainers of the sampler include a discriminator and a prior.

work [13] and Romero (Fig. 2d) from the paper on sampling from continuous distributions [15]. Additionally, we implement tailored architectures Fig. 2e-2h, which are explained in more detail below.

The classical part is implemented with PyTorch and consists of a one convolutional layer and one fully-connected layer. The weights of the convolutional layer (16 feature maps produced by 2×2 kernels) are sampled from a distribution either produced by a classical NN (benchmark) or a PQC. PQC consist of a 4-qubits are implemented in PennyLane. We use `lightning.qubit` noiseless simulator to simulate PQCs.

4.3 Classical benchmark

For the classical benchmark we implement the same adversarial training loop as for the hybrid model. We substitute the PQC with a minimal width and depth NN that reached comparable performance to conventional VI (see Appendix B). This NN consists of a single hidden layer of 8 neurons and 4 input neurons (to match 4 qubit setup).

4.4 Performance metrics

To measure the success of the model, we need to track both, its predictive power as well as the corresponding uncertainty of the predictions. To gauge the predictive performance of the model we track its training and validation accuracy curves and we record additional statistics on the test dataset, such as precision, recall and

F1 score¹. To assess an uncertainty estimation of the model, we track the mean confidence of the final prediction, which allows us to estimate the model’s calibration and confidence error defined as

$$\text{Confidence error} = \text{Mean confidence} - \text{Accuracy}. \quad (12)$$

We strive to create a model that provides a more reliable information on its certainty in the answer. We therefore compare the confidence of both, correctly and incorrectly estimated samples. Additionally, we track the size of an ensemble that contributed towards the final prediction.

5 RESULTS

In this section, we analyse the performance of QCBNN against its classical counterpart. The effects of different PQC architectures on the predictive performance and the uncertainty awareness are tested. We perform this analysis in an iterative fashion as we start (Section 5.1) with the investigation of different architectures from the literature (Fig. 2a-2d) and bring forward hypotheses on which architectural elements might contribute to their successes and failures. We test these hypotheses further (Section 5.2) as we create custom architectures (Fig. 2e-2h) with emphasised properties to see whether these modifications can elevate the performance.

5.1 Architectures from the literature

We first investigate how well a simple setup with a high-order embedding and a varying single calculation layer, that was presented in the literature for similar setups (see Fig. 2a-2d), performs in the QCBNN on the BreastMNIST dataset. We first consider different performance metrics that capture predictive power and uncertainty awareness separately, then we propose a way to fuse these metrics that enables easier analysis.

5.1.1 Predictive performance. Fig. 3a shows training and validation curves and final scores (accuracy, F1 score etc.) of Fig. 2a-2d PQC architectures as well as the performance of the classical benchmark in this setup. All tested architectures struggled to reach the performance of classical benchmark. On the test set all PQCs show more volatile behaviour on recall (model’s ability to spot all cancerous datapoints), which seem to be more sensitive to starting conditions than the benchmark. Interestingly, the PQC architectures *Matic I* and *Matic II* that were used successfully in an analogous setup [10] but performed a different task struggled to learn here. This is visible on both, the highly volatile training and validation curves as well as low final performance. The architectural difference between these PQCs are minor and consist of a mere constraint of rotation, meaning that *Matic I* only has R_x gates while *Matic II* allows to fully parametrized the rotation. Changing the entangling strategy by dropping the last entangling gate, as in *Nikoloska* and *Romero*, creates a behaviour that is closer to the classical benchmark on this dataset. Yet, the training and validation curves are still more volatile and stay below the benchmark. Adding one more layer of (phase) rotation gates, as in *Nikoloska*, seems to make the training curves a little more volatile, however, improves test mean accuracy and recall scores.

¹A harmonic mean of precision and recall, which is a common prediction metric for imbalanced datasets

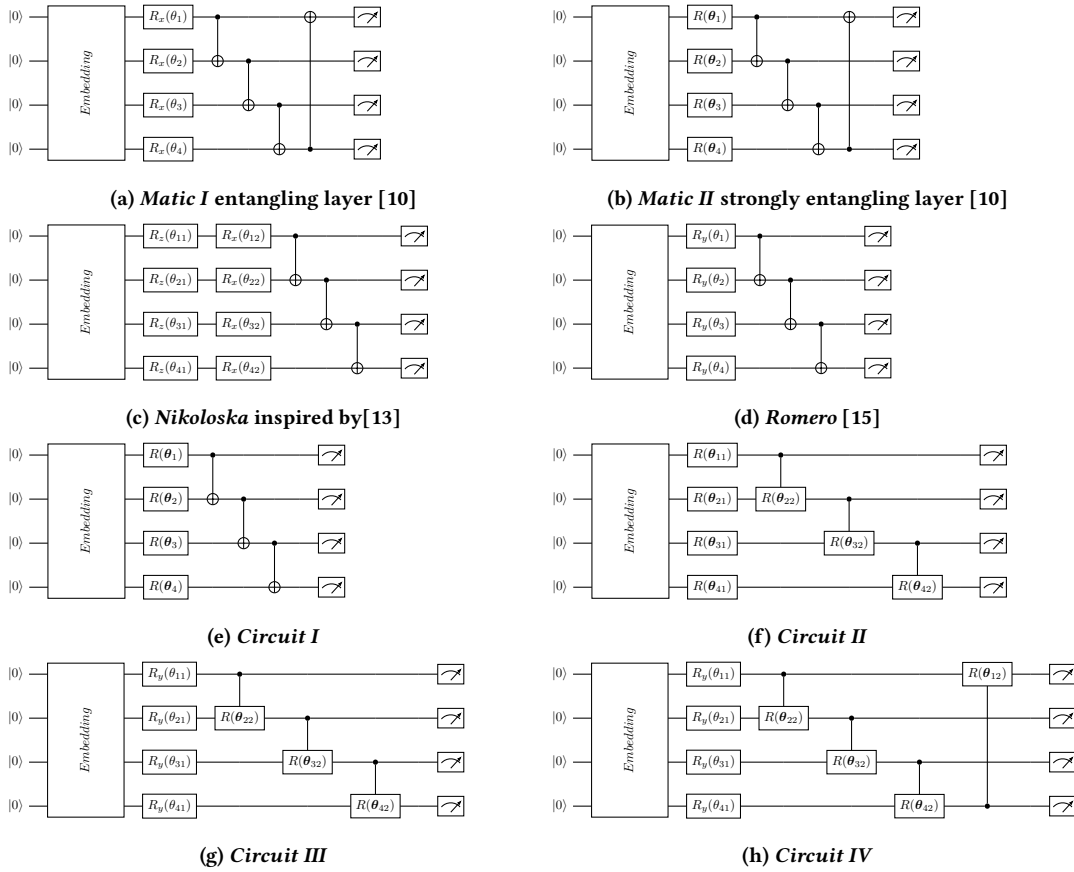



Figure 2: Different PQC architectures, inspired by the literature. All architectures contain an embedding block, trainable block with rotation gate R parametrized by θ followed by measurement of all qubits.


5.1.2 *Uncertainty estimation.* Fig. 3b exposes a different side of the performance metrics that are relevant for BNNs, namely uncertainty metrics. As can be seen from calibration curve in Fig. 3b (right), all models tend to be underconfident for a specific interval (20-80%) of probabilities, while the *Matic II* architecture gravitates towards overconfidence on lower and higher probabilities. The mean confidence error as defined in Eq. (12) (left) and ensemble size plots (middle) show approximate distribution of observations computed by kernel density estimate. The means of these distributions are indicated by crosses and blue errors indicate the direction of a favourable development in comparison to a classical benchmark. We track these values for both correctly and incorrectly identified samples to visualize how confidently wrong the models are. For the mean confidence error, we would like to see a lower (absolute) confidence error for correct samples and a higher confidence error on incorrect samples, which would indicate a higher uncertainty. All models stay in the negative range that indicate underconfidence of all models on the final prediction, which is consistent with the calibration curve readings (right plot). For the ensemble size, we would like to have a higher number of ensemble members to vote for a prediction that was correct than incorrect. Interestingly, classical learners' ensemble members seem to converge to a homogeneous

behaviour and vote the same way no matter what. For quantum learners we see more variety. *Matic I* shows a similarly high confidence error for both correctly and incorrectly identified samples. While *Matic II* improves the mean confidence error, it still differentiates itself substantially from the classical benchmark. Both of these model also show a lesser ensemble size for both correct and incorrect samples. Together, these metrics might hint towards poor learning during the training phase, which was already discussed above. *Nikoloska* showed higher uncertainty for incorrect samples while still exhibiting overall high underconfidence. *Romero's* mean confidence error distribution, on the other hand, overlapped with the one from the classical benchmark, however had a bigger spread.

5.1.3 *Weights distribution.* Now, we take a closer look at how the output of the PQC, which generates the weights distribution, changes as we modify the architectures. Fig. 4 shows that all of PQC's expectation values are biased towards zero. The architectures that had more wide spread distributions tend to have stronger performances in terms of metrics considered above. Keeping this in mind, we proceed to tweaking the architectures in the following sections to see if we can force the weights distributions to take a more expressive form and see whether this maps well to improved performance.



(a) Training (left), validation (center) curves and final test performance (right). Lines of the curves (left and center) indicate mean performance and transparent sleeves indicate standard deviation over four random seeds. The box plots of the final metrics (right) show the quartiles of performances generated by the seeds, while whiskers indicate the spread of the distribution.



(b) Mean confidence error (left), ensemble size (center) and calibration (right) on test set. Kernel density estimate plots (left and center) show estimated probability distribution of final uncertainty estimation metrics over four random seeds. The X-markers show the means of the distributions (blue marker shows mean of the classical benchmark), while the arrows show the direction of the desired development.

Figure 3: Overview over predictive performance metrics and uncertainty-related metrics for Fig. 2a-2d PQCs.




Figure 4: Weights distribution density estimated from 100 inference passes through the model and consists of all samples across all convolutional filters combined.

5.1.4 *Choosing architecture.* To find the best performing PQC we need to consider multiple behavioural parameters at once. To ease the analysis, we fuse the uncertainty metrics and compare it to the final accuracy. Given that quantum models show a bigger discrepancy between the ensemble size of correct and incorrect predictions, we weight the mean uncertainty with the ensemble size to acquire a more reliable uncertainty estimation. It is our ambition to find models that possess more dependable point-wise uncertainty estimation capabilities without sacrificing predictive accuracy. We therefore consider whether a PQC provides for a bigger difference in confidence between correctly and incorrectly identified samples. This difference is computed as follows:

$$\begin{aligned} \text{Difference} = & \text{Confidence}_{\text{correct}} \times \text{Ensemble fraction}_{\text{correct}} \\ & - \text{Confidence}_{\text{incorrect}} \times \text{Ensemble fraction}_{\text{incorrect}} \end{aligned} \quad (13)$$

Fig. 5 shows the estimated distribution of these differences versus accuracy of prediction for all PQC architectures considered above. This graph indicates that the PQCs that showed the best

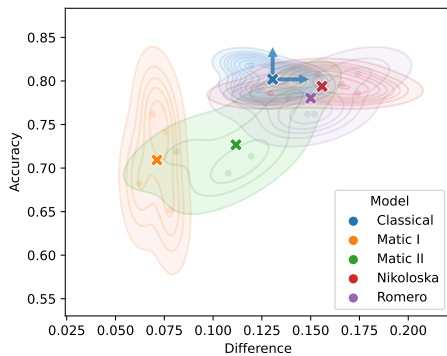


Figure 5: Accuracy on the test set and certainty level difference between correctly and incorrectly identified samples weighted by ensemble size (see Eq. (13)).

performances above (*Nikoloska* and *Romero*) are showing on average a slightly higher confidence difference between correctly and incorrectly identified cases at the expense of a drop in accuracy compared to the benchmark. The distributions of these architectures are significantly overlapping, but *Nikoloska* seems to have a more stable performance in terms of accuracy and slightly higher average difference. This architecture, however, also showed a higher mean confidence error (meaning that the model was less confident in its predictions on average) as can be seen Fig. 3b, which is not reflected in the difference metric.

From the above analysis we conclude that *Romero* architecture had one of the strongest predictive performances on all training, validation and test datasets, the lowest average mean confidence error for the correct prediction, one of the largest confidence gap between correct and incorrect samples, and the most wide spread of the weights distribution generated.

5.2 Custom architectures


PQC architectures that were previously proposed in the literature, unfortunately, failed to deliver favourable results in comparison to classical benchmark for the task at hand as was discussed above. Moreover, the distribution of weights that quantum samplers generated had a strong bias towards 0., which challenges the common argument for using QC for probabilistic ML such as representing distributions that are hard to model classically. Here, we built up upon the insights from the previous section in an attempt to create stronger architectures for our problem. Allowing the architectures a full instead of partial control, as in going from *Matic I* to *Matic II*, boosted the predictive capability. Therefore, we perform the same trick with the best performing architecture from the literature *Romero* and create *Circuit I* (Fig. 2e). Another acquired insight from the previous study is that the change in the entanglement layer had the most profound impact on the PQC’s behaviour (moving from *Matic I* and *Matic II* to *Nikoloska* and *Romero*). It was beneficial for both, predictive performance and uncertainty awareness. Therefore, we change the entangling layer from CNOT to CR gates for *Romero* and *Circuit I* and create *Circuit III* (Fig. 2g) and *Circuit II* (Fig. 2f) respectively to allow the model to tune its own entangling

during training. Additionally, we include *Circuit IV* where we add an entangling gate between first and last qubit to test whether this entanglement strategy is generally not fitting for this task or not.


5.2.1 Predictive performance. Fig. 6a shows the performance of different custom architectures (Fig. 2e - Fig. 2h). Unexpectedly, providing more freedom to the PQC to tweak its rotation with *Circuit I* hindered the training in comparison to *Romero* (see Fig. 3a), which seems to indicate that there is a possible saturation point beyond which a free choice of rotation is no longer boosting the model’s capacity. Proving more freedom for the entangling layers, on the other hand, lead to the expected boost: *Circuit II* provided an improved performance over *Circuit I* and previously tested PQC in Fig. 3a. Its performance on training, validation and test set is, however, still on average a bit lower than of the classical benchmark. Restricting the rotation layer the same way as in *Romero* (Fig. 2d) for architectures *Circuit III* and *Circuit IV* provided enough capacity to go slightly beyond the classical benchmark on the training dataset. Here, we see barely any difference between two entangling strategies in *Circuit III* and *Circuit IV*. Interestingly, even though there is a clear distinction of the training curves on the training data, the ones on validation data show almost no discrepancy. On the test dataset, *Circuit III* managed to slightly outperform the classical benchmark in terms of the average accuracy, but the spread of accuracies is high. Looking at other metrics, we can see that quantum models tend to provide a higher precision (high quality of solution), but suffer from wide spread in recall, meaning that, dependent on initial conditions, our model managed to either over- or underperform in comparison to the classical benchmark.

5.2.2 Uncertainty estimation. Fig. 6b provides an overview of different uncertainty measures of the model. Here again, as in Fig. 3b, we see barely any difference in calibration curves of the models. The mean confidence error distributions, however, tell a slightly different story: *Circuit I* that had trainability issues, as discussed above, also shows lower mean confidence error on correct samples and higher error on incorrect samples. *Circuit II* on average shows a slightly higher uncertainty on incorrectly identified samples, without rising the uncertainty of the correctly identified samples. This analysis is valid on average, but given the spread of the error distribution, each single model might show a bit different behaviour. Constraining the rotation layer of the previously discussed model to *Circuit III* leads to lesser spread of the error distribution and lesser mean confidence error for both correctly and incorrectly identified samples. Changing the entanglement strategy from *Circuit III* to *Circuit IV* leads to a smaller on average mean confidence error for both correctly and incorrectly classified instances. In terms of ensemble size we see a favourable development for all quantum models as the correct samples receive more votes in favour, while for incorrect samples ensemble agree less.

5.2.3 Weights distribution. Now, we analyse the distributions that were generated by the PQCs. Fig. 7 provides an overview of the effects of different architectures on the actual output of the PQCs. *Circuit I*’s output that had the most constrained entanglement regime is still biased towards 0. (the same as for architectures analysed previously), however, changing the entanglement layer allowed the distributions to be wider and take more expressive forms. From



(a) Training (left), validation (center) curves and final test performance (right). Lines of the curves (left and center) indicate mean performance and transparent sleeves indicate standard deviation over four random seeds. The box plots of the final metrics (right) show the quartiles of performances generated by the seeds, while whiskers indicate the spread of the distribution.



(b) Mean confidence error (left), ensemble size (center) and calibration (right) on test set. Kernel density estimate plots (left and center) show estimated probability distribution of final uncertainty estimation metrics over four random seeds. The X-markers show the means of the distributions (blue marker shows mean of the classical benchmark), while the arrows show the direction of the desired development.

Figure 6: Overview over predictive performance metrics and uncertainty-related metrics for Fig. 2e-2h PQCs.

this graph we see that the models that performed at least the same as a classical benchmark also have at maximum the same density level, which aligns with a hypothesis from above that BNN favour generators that provide diverse outputs.

5.2.4 *Choosing architecture.* Here, we perform the same analysis as was done in Section 5.1.4 for tweaked PQC architectures Fig. 2e - 2h. Fig. 8 shows the distributions of differences and accuracies across different architectures. All quantum architectures show on average a slightly higher difference between correctly and incorrectly identified samples, which comes at an expense of a lower average accuracy. The lowest performing PQC *Circuit I* shows a wide spread of the difference distribution. Relaxing the entanglement layers in *Circuit II* leads to a much more narrow distribution, which reaches the differences at least as good as the classical benchmark. The average predictive performance is, however, slightly lower. Changing the entanglement strategy as in *Circuit III* leads to a higher difference, which overlaps with classical benchmark only a little. The predictive performance of this model greatly depends on initial conditions as discussed above. Adding one additional




Figure 7: Weights distribution density estimated from 100 inference passes through the model and consists of all samples across all convolutional filters combined.

entanglement gates as in *Circuit IV* slightly lowers the average performance and increases the spread of the distribution.




Figure 8: Accuracy on the test set and certainty level difference between correctly and incorrectly identified samples, weighted by ensemble size (see Eq. (13)).

Based on results from above, we draw the conclusion that PQC architecture *Circuit III* had the strongest performance compared to *Circuit I, II* and *IV*, as well as PQCs *Matic I, Matic II, Nikoloska* and *Romero* from the previous section. *Circuit III* also slightly exceeded the classical benchmark in predictive performance and uncertainty estimation. To test whether the performance of *Circuit III* has saturated we performed additional testing in Appendix A.

6 DISCUSSION

The results described in the previous section imply that by introducing learnable parameters into the entangling layer allows to achieve better performances for this task, which is not a common practice that we see across the literature. It correlates with weights distributions having a wider spread, which seem to be more favoured by the setup of a QCBNN. On the contrary, in the setup of QCCNN [10] the PQCs *Matic I* outperformed *Matic II*, while in our setup these PQCs are the worst performing ones and fall below the classical benchmark. This could indicate that the tasks of computing convolutional operation (as in QCCNN) and sampling from weights distribution (as in QCBNN) favour contrasting behavioural patterns of the models.

An interesting further research direction is to provide more theoretical basis that would allow to construct tailored PQCs in a less heuristic fashion. One way to approach this is to investigate the connection of these circuits to the coefficients of Fourier series that can be fitted to the periodic output of the PQCs as proposed in [18]. In that work, the authors showed that the embedding of a circuit controls the number of frequencies in the series and thus the expressivity of the model (the more frequencies are available, the more complex function one can represent with the models). However, it is speculated that the structure of the computation layer effects the coefficients of these series, but this idea was not developed further. In our work, we have shown that some architectural decisions of PQCs resulted in more wide distributions of the weights, and it might prove to be beneficial to investigate this point more thoroughly by looking into the connection of these decisions to the coefficients. A related further research direction is a connection of the PQCs to classical surrogates as introduced in [17].

Provided that the structure of selected PQCs fulfills the criteria, it is possible to build a classical surrogate model that could serve as a more reliable classical benchmark as well as ease deployability of these models. Allowing too much freedom in rotational parameters is, however, detrimental to the performance (e.g. *Circuit II* to *Circuit III*). This could be attributed to more complex loss landscapes that hinder the optimizer to find optima. However, a deeper dive into the origins of the problem might shed some light onto why some PQC architectures are more successful than others.

One thing that becomes apparent through out experiments described above is that the training curves of the PQC in Fig. 3a and Fig. 6a do not necessarily map well to the accuracies of these architectures on the test dataset, which could be attributed to a rather small test dataset. In future work, we are planning to perform further experiments to test how these results map to bigger datasets. Furthermore, in this work, we do not distinguish between epistemic and aleatoric uncertainties, like it was done in e.g. [12], which is an interesting additional metric to consider in the future works.

In a pursuit of allowing QCBNN model to represent continuous weights and hence learn more complex datasets we modified Born machine according to [15]. This entailed that the source of stochasticity was shifted from quantum to classical devices. A further interesting direction of research could be to change the way we sample from the generator to allow the PQC to remain a Born machine and yet keep the possibility for weight to remain continuous. One idea to do this would be to implement something akin Monte Carlo Dropout methods, where PQC can be responsible for selecting the weights that are dropped out during this iteration. There are a few challenges to overcome with this idea, e.g. dropout is often implemented as a Bernoulli distribution, so what type of PQC architecture could be beneficial here and whether or not it should have learnable parameters remains to be investigated.

7 CONCLUSION

In this work, we introduce a Quantum-Classical Bayesian Neural Network (QCBNN) with continuous weights for uncertainty-aware classification of breast ultrasound scans (BreastMNIST). This model is designed as a fusion of a previously introduced models [10, 13, 15] and it is evaluated from the stand points of its predictive performance and uncertainty awareness. Many behavioral aspects of the models are considered that are captured by a variety of metrics. We methodically test multiple PQC architectures in this scenario, including the ones that has been successfully used in a similar setup before (Fig. 2a-2d). We determined that certain architectural features, such as trainable entanglement layers and rotation layers with less parameters, allow for better learning on this dataset. We test these hypotheses by building custom PQCs for this task (Fig. 2e-2h) that boost the performance and even allow to slightly outperform the classical benchmark in terms of uncertainty awareness. These results provide interesting empirical insights for more informed PQC designs in the future.

ACKNOWLEDGMENTS

This research is supported by the Bavarian Ministry of Economic Affairs, Regional Development and Energy with funds from the

Hightech Agenda Bayern. We thank Marc Machaczek for his role in the preparatory stage of the project.

A DIFFERENT NUMBER OF LAYERS FOR CIRCUIT III

Here, we investigate whether the performance of best performing *Circuit III* presented in Section 5.2 can benefit from deepening its architecture. The layers can be added by simple repeating the computation layer, or adding both embedding and computation layers together to create a re-uploading model [14].

Fig. 9 shows predictive performance of *Circuit III* (Fig. 2g) with one layer (*L1*, the same setup as in Section 5.2), with two layers (*L2*) and two layers in re-uploading setup (*L2 Re*). Increasing the depth from *L1* to *L2* seems to affect the trainability as the average training curve falls below one layer setup. Adding additional embedding layer before *L2* (*L2 Re*) seems to hinder the trainability even further, as performances on training, validation and test datasets are impaired.

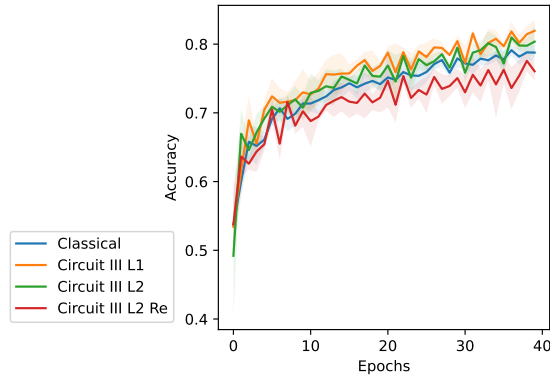


Figure 9: Training curves for Fig. 2g PQC with varying depth. Lines of training and validation curve indicate mean performance and transparent sleeves indicate standard distribution over four random seeds.

Next, we consider test accuracy and difference confidence levels between correctly and incorrectly identified samples (as in Eq. (13)). Fig. 10 displays these metrics for different number of layers. One layer *Circuit III L1* PQC shows higher confidence gap but quite volatile predictive performance as final accuracy between four random seeds differ a lot. Adding another computation layer (*L2*) stabilizes the predictive performance to the levels similar to the classical benchmark, but destabilizes the uncertainty estimation capacity.

From the above analysis we conclude that increasing the amount of layers two *L2* stabilizes the predictive performance on testing set while proving a larger confidence gap on average. The average performance on training dataset on the other hand dropped a little, which could hint towards drop in trainability.

B VARIATIONAL INFERENCE

Here, we show the comparison between a standard approach to training classical BNNs namely the Variational Inference (VI) and the Adversarial training (Adv) that serve as a classical benchmark.




Figure 10: Prediction performance as measured by accuracy on the test set and uncertainty estimation difference between correctly and incorrectly identified samples that is weighted by ensemble size (see Eq. (13)).

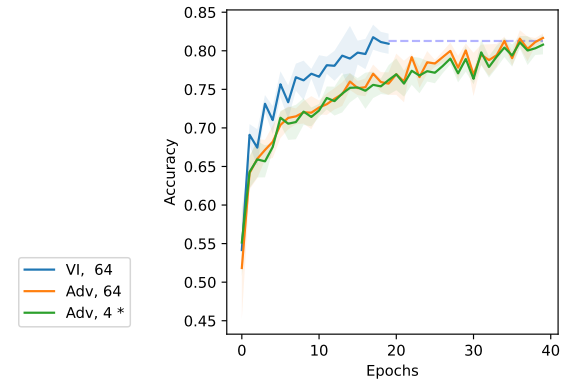


Figure 11: Variational Inference (VI) vs Adversarial learning loop (Adv) training curves of classical BNNs. The numbers next to method name indicate the size of the injected noise vector.

Fig. 11 shows that Adv achieves the same accuracy as VI method even with substantially smaller injected noise vector, however, it take more epochs to do so.

REFERENCES

- [1] Amira Abbas, David Sutter, Christa Zoufal, Aurelien Lucchi, Alessio Figalli, and Stefan Woerner. 2021. The power of quantum neural networks. *Nature Computational Science* 1, 6 (June 2021), 403–409. <https://doi.org/10.1038/s43588-021-00084-1>
- [2] Marcello Benedetti, Brian Coyle, Mattia Fiorentini, Michael Lubasch, and Matthias Rosenkranz. 2021. Variational Inference with a Quantum Computer. *Physical Review Applied* 16, 4 (oct 2021). <https://doi.org/10.1103/physrevapplied.16.044057>
- [3] Noah Berner, Vincent Fortuin, and Jonas Landman. 2021. Quantum Bayesian Neural Networks. arXiv:2107.09599 [quant-ph]
- [4] Brian Coyle, Maxwell Henderson, Justin Chan Jin Le, Niraj Kumar, Marco Painsi, and Elham Kashefi. 2020. Quantum versus Classical Generative Modelling in Finance. arXiv:2008.00691 [quant-ph]
- [5] Brian Coyle, Daniel Mills, Vincent Danos, and Elham Kashefi. 2020. The Born supremacy: quantum advantage and training of an Ising Born machine. *npj Quantum Information* 6, 1 (jul 2020). <https://doi.org/10.1038/s41534-020-00288-9>
- [6] Kaitlin Gili, Mohamed Hibat-Allah, Marta Mauri, Chris Ballance, and Alejandro Perdomo-Ortiz. 2022. Do Quantum Circuit Born Machines Generalize? <https://arxiv.org/abs/2205.12345>

- [//doi.org/10.48550/ARXIV.2207.13645](https://doi.org/10.48550/ARXIV.2207.13645)
- [7] Hsin-Yuan Huang, Michael Broughton, Masoud Mohseni, Ryan Babbush, Sergio Boixo, Hartmut Neven, and Jarrod R. McClean. 2021. Power of data in quantum machine learning. *Nature Communications* 12, 1 (May 2021). <https://doi.org/10.1038/s41467-021-22539-9>
- [8] Prateek Jain and Srinjoy Ganguly. 2022. Hybrid Quantum Generative Adversarial Networks for Molecular Simulation and Drug Discovery. arXiv:2212.07826 [quant-ph]
- [9] Laurent Valentin Jospin, Hamid Laga, Farid Boussaid, Wray Buntine, and Mohammed Bennamoun. 2022. Hands-On Bayesian Neural Networks—A Tutorial for Deep Learning Users. *IEEE Computational Intelligence Magazine* 17, 2 (may 2022), 29–48. <https://doi.org/10.1109/mci.2022.3155327>
- [10] Andrea Matic, Maureen Monnet, Jeanette Miriam Lorenz, Balthasar Schachtner, and Thomas Messerer. 2022. Quantum-classical convolutional neural networks in radiological image classification. arXiv:2204.12390 [quant-ph]
- [11] Maureen Monnet, Hanady Gebran, Andrea Matic-Flierl, Florian Kiwit, Balthasar Schachtner, Amine Bentellis, and Jeanette Miriam Lorenz. 2023. Pooling techniques in hybrid quantum-classical convolutional neural networks. arXiv:2305.05603 [quant-ph]
- [12] Nam Nguyen and Kwang-Cheng Chen. 2022. Bayesian Quantum Neural Networks. *IEEE Access* 10 (2022), 54110–54122. <https://doi.org/10.1109/ACCESS.2022.3168675>
- [13] Ivana Nikoloska and Osvaldo Simeone. 2022. Quantum-Aided Meta-Learning for Bayesian Binary Neural Networks via Born Machines. <https://doi.org/10.48550/ARXIV.2203.17089>
- [14] Adrián Pérez-Salinas, Alba Cervera-Lierta, Elies Gil-Fuster, and José I. Latorre. 2020. Data re-uploading for a universal quantum classifier. *Quantum* 4 (Feb. 2020), 226. <https://doi.org/10.22331/q-2020-02-06-226>
- [15] Jonathan Romero and Alan Aspuru-Guzik. 2019. Variational quantum generators: Generative adversarial quantum machine learning for continuous distributions. <https://doi.org/10.48550/ARXIV.1901.00848>
- [16] Manuel S. Rudolph, Ntwali Bashige Toussaint, Amara Katarbarwa, Sonika Johri, Borja Peropadre, and Alejandro Perdomo-Ortiz. 2022. Generation of High-Resolution Handwritten Digits with an Ion-Trap Quantum Computer. arXiv:2012.03924 [quant-ph]
- [17] Franz J. Schreiber, Jens Eisert, and Johannes Jakob Meyer. 2023. Classical Surrogates for Quantum Learning Models. *Physical Review Letters* 131, 10 (Sept. 2023). <https://doi.org/10.1103/physrevlett.131.100803>
- [18] Maria Schuld, Ryan Sweke, and Johannes Jakob Meyer. 2021. Effect of data encoding on the expressive power of variational quantum-machine-learning models. *Physical Review A* 103, 3 (March 2021). <https://doi.org/10.1103/physreva.103.032430>
- [19] Jiancheng Yang, Rui Shi, Donglai Wei, Zequan Liu, Lin Zhao, Bilian Ke, Hanspeter Pfister, and Bingbing Ni. 2023. MedMNIST v2-A large-scale lightweight benchmark for 2D and 3D biomedical image classification. *Scientific Data* 10, 1 (2023), 41.
- [20] Zhikuan Zhao, Alejandro Pozas-Kerstjens, Patrick Rebstrodt, and Peter Wittek. 2019. Bayesian deep learning on a quantum computer. *Quantum Machine Intelligence* 1, 1–2 (May 2019), 41–51. <https://doi.org/10.1007/s42484-019-00004-7>
- [21] Christa Zoufal, Aurélien Lucchi, and Stefan Woerner. 2019. Quantum Generative Adversarial Networks for learning and loading random distributions. *npj Quantum Information* 5, 1 (Nov. 2019). <https://doi.org/10.1038/s41534-019-0223-2>

Received 15 April 2024; revised 4 June 2024