

QMetric: Benchmarking Quantum Neural Networks Across Circuits, Features, and Training Dimensions

Silvie Illéssová^{1,*}, Tomasz Rybotycki^{2,3,4,†} and Martin Beseda^{5,†}

¹IT4Innovations National Supercomputing Center, Studentská 6231/1B, 708 00 Ostrava, Czech Republic

²Systems Research Institute, Polish Academy of Sciences, ul. Newelska 6, 01-447 Warszawa, Poland

³Nicolaus Copernicus Astronomical Center, Polish Academy of Sciences, ul. Bartycka 18, 00-716 Warsaw, Poland

⁴Center of Excellence in Artificial Intelligence, AGH University, Aleje Mickiewicza 30, 30-059 Cracow, Poland

⁵Department of Information Engineering, Computer Science and Mathematics, University of L'Aquila, via Vetoio, I-67010 Coppito-L'Aquila, Italy

Abstract

As hybrid quantum-classical models gain traction in machine learning, there is a growing need for tools that assess their effectiveness beyond raw accuracy. We present **QMetric**, a Python package offering a suite of interpretable metrics to evaluate quantum circuit expressibility, feature representations, and training dynamics. QMetric quantifies key aspects, including circuit fidelity, entanglement entropy, barren plateau risk, and training stability. The package integrates with Qiskit and PyTorch, and is demonstrated via a case study on binary MNIST classification comparing classical and quantum-enhanced models. Code, plots, and a reproducible environment are available on GitHub.

Keywords

Quantum Computing, Machine Learning, Hybrid Quantum-Classical Model, Quantum Machine Learning, Evaluation Metrics, Benchmarking

1. Introduction

Hybrid quantum-classical neural networks (QNNs) [1, 2, 3] are playing a central role in the development of algorithms for near-term quantum devices. By embedding parameterized quantum circuits within classical training loops, these architectures aim to leverage quantum resources such as entanglement and superposition while maintaining trainability through well-established classical optimizers. This hybrid structure has enabled a broad spectrum of quantum machine learning (QML) [4, 5] models to flourish across domains including classification [6, 7], generative modeling [8, 9], optimization [10, 11, 12], benchmarking [13, 14], medicine [15, 16, 17, 18], and quantum chemistry [19, 20, 21].

Importantly, many canonical variational algorithms—originally developed for quantum simulation—can be reframed as learning architectures. The *Variational Quantum Eigensolver* (VQE) [22, 23, 24, 25] exemplifies this duality. Therein, a parameterized quantum ansatz is trained to minimize an energy objective, similarly to a neural network minimizing a loss function. Over time, VQE has evolved into a family of learning-based formulations, including *State-Averaged Orbital-Optimized VQE* [26, 27], *ADAPT-VQE* [28], and *Subspace-Search VQE* [29], each introducing novel strategies for parameterization, target state selection, and optimization flow.

Beyond simulation, hybrid QNNs are widely applied in supervised learning, typically in classification or regression tasks [30]. Here, quantum circuits are used to encode classical data (via feature maps [31, 32]), process it through a variational ansatz, and output probabilities or decision boundaries. Such architectures are used in *Quantum Neural Networks* [33], *Quantum Support Vector Machines* [34], and more recent paradigms like *Quantum Kitchen Sinks* [35] and *Quantum Feature Spaces* [36]. In

QualITA 2025: The Fourth Conference on System and Service Quality, June 25 and 27, 2025, Catania, Italy

*Corresponding author.

†These authors contributed equally.

✉ illesova.silvie.scholar@gmail.com (S. Illéssová); tryb@camk.edu.pl (T. Rybotycki); martin.beseda@univaq.it (M. Beseda)

ORCID 0009-0002-5231-3714 (S. Illéssová); 0000-0003-2493-0459 (T. Rybotycki); 0000-0001-5792-2872 (M. Beseda)



© 2025 Copyright for this paper by its authors. Use permitted under Creative Commons License Attribution 4.0 International (CC BY 4.0).

unsupervised learning, models such as *Quantum Circuit Born Machines* [37], and *Quantum Autoencoders* [38] extend the reach of QML into generative and latent-variable modeling.

Despite the growing variety and sophistication of QML models, there remains a lack of principled, interpretable, and reproducible tools for evaluating their behavior. Traditional ML diagnostics—accuracy, F1-score, or validation loss—do not capture key quantum characteristics such as circuit expressibility, entanglement structure, barren plateaus, or the sensitivity of quantum feature maps. Without such metrics, model design becomes largely heuristic, and comparisons between quantum and classical architectures are often inconclusive or misleading.

To bridge this gap, we introduce **QMetric**, a modular and extensible Python framework for evaluating hybrid quantum-classical models. QMetric computes interpretable scalar metrics across three complementary dimensions (i) the structure and expressiveness of quantum circuits; (ii) the geometry and compression of quantum feature spaces; and (iii) the stability, efficiency, and gradient flow during training. These tools allow researchers to diagnose bottlenecks, compare architectures, and validate empirical claims beyond raw accuracy.

Our package integrates with Qiskit ¹ and PyTorch ², which we demonstrate through a binary classification example on MNIST digits. Therein, we compare a classical neural network with a hybrid QNN. All code, plots, and environment files are publicly available for reproducibility and further experimentation.

2. Software Specifications

All experiments were conducted using the `qmetric-env` Conda ³ environment ⁴, configured for hybrid quantum-classical machine learning. The system exploits GPU-accelerated libraries, supports Qiskit primitives of version V1, and integrates PyTorch and scikit-learn⁵ for classical model components and preprocessing.

The environment is based on Python 3.10.13 with key libraries and versions listed in Table 1. Qiskit version 1.4.3 was used in conjunction with Qiskit Aer 0.17.0 and Qiskit Machine Learning 0.8.2. PyTorch version 2.7.0 and CUDA 12.9 toolchain were used for classical and hybrid model execution. Principal component analysis and classical baseline training relied on scikit-learn version 1.6.1.

Table 1
Key Software Components

Library	Version
Python	3.10.13
Qiskit	1.4.3
Qiskit Aer	0.17.0
Qiskit Machine Learning	0.8.2
PyTorch	2.7.0
CUDA Toolkit	12.9
cuDNN	9.10.1.4
scikit-learn	1.6.1
NumPy	2.2.6
Matplotlib	3.10.3
SymPy	1.14.0

The training experiments were run on a Linux system using conda’s `qmetric-env` environment. Hardware acceleration via CUDA and cuDNN was enabled to support efficient execution of neural net-

¹<https://www.ibm.com/quantum/qiskit>

²<https://pytorch.org/>

³<https://anaconda.org/anaconda/conda>

⁴<https://gitlab.com/illesova.silvie.scholar/qmetric/-/blob/main/environment.yml>

⁵<https://scikit-learn.org/stable/>

work operations and gradient computation. Quantum simulations were executed using AerSimulator in statevector mode.

Note that the Qiskit primitives interface used in the hybrid model, `EstimatorQNN`⁶, is marked as deprecated in favor of V2 primitives. Future implementations of QMetric should migrate to the `StatevectorEstimator`⁷ to ensure compatibility with upcoming Qiskit releases.

3. Metrics Categories

QMetric organizes its metrics into three complementary categories—quantum circuit behavior, quantum feature space, and training dynamics—that together provide a comprehensive profile of a hybrid model’s expressiveness, learnability, and robustness. These categories are summarized in Tab. 2.

3.1. Quantum Circuit Metrics

As the computational core of hybrid models, quantum circuits influence representational capacity and noise resilience. QMetric evaluates circuit quality through metrics such as *Quantum Circuit Expressibility*, which measures the diversity of quantum states produced under random parameters, and *Quantum Circuit Fidelity*, estimating robustness to noise via state overlap.

To characterize circuit structure, the *Quantum Locality Ratio* captures the balance between local and entangling gates. Entanglement-based metrics include *Effective Entanglement Entropy* and *Quantum Mutual Information*, which quantify intra-circuit quantum correlations. These metrics are useful when tuning ansätze for VQE, QAOA, or classification tasks, where poor ansatz expressibility or excessive entanglement can hinder learning.

3.1.1. Quantum Circuit Expressibility

Quantum Circuit Expressibility (QCE) [39] quantifies a circuit’s ability to generate a diverse set of quantum states across the Hilbert space. It measures how closely the distribution of states produced by the parameterized circuit approximates the uniform (Haar) distribution [40]. High expressibility corresponds to broader state coverage in Hilbert space and is conceptually linked to the Fubini-Study distance [41]. It can be quantitatively related to the Kullback–Leibler divergence [42] between the circuit’s output distribution and the Haar distribution. QCE implies that the circuit can reach a wide variety of states, which is crucial for representing complex functions in quantum machine learning and variational algorithms.

Formally, QCE is defined via the pairwise fidelity of randomly generated state vectors,

$$\text{QCE} = 1 - \frac{1}{N(N-1)} \sum_{i < j} |\langle \psi_i | \psi_j \rangle|^2, \quad (1)$$

where N is the number of randomly sampled parameter sets used to generate the corresponding quantum states, $\{\psi_i\}_{i=1}^N$ are the quantum states obtained by randomly sampling parameters from the specified ranges and applying them to the circuit. This expression captures the average overlap between states. Lower overlap corresponds to greater expressibility. The QCE score lies in the range $[0, 1]$, with values closer to 1 indicating higher expressiveness.

In practice, QCE helps identify whether a variational circuit is too shallow (low expressibility) or overly complex (potentially prone to barren plateaus). A well-designed circuit should maintain a high QCE while preserving trainability and manageable entanglement. QMetric implements QCE by sampling multiple parameter sets, evaluating state vector overlaps, and averaging across all pairwise fidelities, making it an efficient diagnostic for early-stage ansatz evaluation.

⁶https://qiskit-community.github.io/qiskit-machine-learning/stubs/qiskit_machine_learning_neural_networks.EstimatorQNN.html

⁷<https://quantum.cloud.ibm.com/docs/en/api/qiskit/qiskit.primitives.SamplerPubResult>

3.1.2. Quantum Circuit Fidelity

Quantum Circuit Fidelity (QCF) [43] quantifies the robustness of a quantum circuit to noise by measuring how closely the output of the noisy circuit resembles that of the ideal (noise-free) version. Fidelity serves as a key metric for assessing noise resilience in near-term quantum devices, where decoherence, gate errors, and readout noise can significantly degrade quantum state quality.

Mathematically, QCF is defined as the fidelity between two quantum states,

$$F(\rho, \sigma) = \left(\text{Tr} \sqrt{\sqrt{\rho} \sigma \sqrt{\rho}} \right)^2, \quad (2)$$

where ρ is the density matrix of the ideal output state and σ represents the output state under a specified noise model. In the special case of pure states (as in most simulation scenarios), the fidelity simplifies to the squared absolute value of the inner product between the ideal and noisy state vectors.

In QMetric, QCF is computed by simulating both the ideal and noisy execution of a circuit using Qiskit's statevector simulator and a user-defined noise model. The resulting fidelity score ranges from 0 to 1, with higher values indicating stronger fidelity. QCF is especially useful when benchmarking circuits across different hardware targets or when optimizing ansatz designs for noisy intermediate-scale quantum (NISQ) devices [44].

3.1.3. Quantum Locality Ratio

Quantum Locality Ratio (QLR) [45] quantifies the proportion of single-qubit operations relative to the total number of gates in a quantum circuit. This metric captures the locality of interactions, offering insight into how much a circuit relies on entangling operations. A high QLR implies that the circuit uses mostly local, single-qubit gates, whereas a low value suggests strong reliance on multi-qubit entanglement.

Formally, QLR is defined as

$$\text{QLR} = \frac{N_{1-q}}{N_{\text{total}}}, \quad (3)$$

where N_{1-q} denotes the number of gates acting on a single qubit and N_{total} is the total number of gates in the circuit.

In QMetric, this ratio is computed by iterating over the circuit's gate operations and counting how many act on one qubit. QLR helps to assess the tradeoff between locality and entanglement. It returns the ratio of single to multi-qubit gates, providing a fast and interpretable structural descriptor. QLR is particularly useful during ansatz design, where excessive entanglement can introduce barren plateaus or hardware noise sensitivity.

3.1.4. Effective Entanglement Entropy

Effective Entanglement Entropy (EEE) [46] evaluates the degree of quantum entanglement between a subsystem of qubits and the rest of the circuit. It is based on the von Neumann entropy of the reduced density matrix of the selected subsystem, capturing how mixed its state becomes due to entanglement with its complement.

The metric is defined as

$$S(\rho_A) = -\text{Tr}(\rho_A \log \rho_A), \quad (4)$$

where ρ_A is the reduced density matrix obtained by tracing out all qubits missing from the chosen subsystem.

QMetric computes EEE by generating a state vector from the circuit, selecting a target subset of qubits, performing a partial trace, and evaluating the entropy. This metric is useful in tasks like entanglement scaling analysis, where understanding subsystem correlations is essential for tuning circuit depth and topology.

3.1.5. Quantum Mutual Information

Quantum Mutual Information (QMI) [47] measures the total correlations—both classical and quantum—between two disjoint subsets of qubits in a quantum circuit. It extends the concept of mutual information to the quantum domain, revealing how strongly two regions of a circuit are statistically linked.

The metric is computed as

$$I(A : B) = S(\rho_A) + S(\rho_B) - S(\rho_{AB}), \quad (5)$$

where ρ_A , ρ_B , and ρ_{AB} are the reduced density matrices of subsystems A , B , and their union, respectively.

In QMetric, QMI is computed by preparing a full state vector, currently via analytical simulation, computing partial traces for each subsystem and their union, and evaluating the entropies. This metric is instrumental for analyzing modular architectures, verifying disentanglement, or diagnosing undesired correlations in VQE, QAOA, or classification-oriented quantum circuits [48].

3.2. Quantum Feature Space Metrics

When encoding classical data into Hilbert space, the geometry of the resulting feature space directly affects model performance. QMetric provides the *Feature Map Compression Ratio* (FMCR), assessing how efficiently classical data are compressed via PCA, and the *Effective Dimension* (EDQFS), which reflects variance spread in the quantum feature space.

The *Quantum Layer Activation Diversity* (QLAD) and *Quantum Output Sensitivity* (QOS) evaluate output variability and robustness to perturbations. Low QLAD and high QOS signal collapsed or brittle encodings. These metrics are critical in Parametrized Quantum Circuit (PQC)-based classifiers, quantum kernel methods, and other models relying on quantum feature geometry.

3.2.1. Feature Map Compression Ratio

Feature Map Compression Ratio (FMCR) [49] quantifies how efficiently a quantum feature map compresses the input data. It compares the original classical dimensionality with the number of principal components needed to capture most of the variance in the quantum-transformed space. A high FMCR indicates strong compression, meaning fewer effective dimensions are required to retain the majority of the encoded information.

Formally, FMCR is defined as

$$\text{FMCR} = \frac{d_{\text{in}}}{d_{\text{eff}}}, \quad (6)$$

where d_{in} is the dimensionality of the classical input and d_{eff} is the number of principal components explaining 95% of the variance in the quantum feature space.

QMetric implements FMCR by applying PCA to the quantum-transformed dataset, calculating the cumulative explained variance, and identifying the number of components required to exceed the 95% threshold. This metric is especially relevant when assessing whether a feature map leads to redundancy or useful abstraction..

3.2.2. Effective Dimension of Quantum Feature Space

Effective Dimension of Quantum Feature Space (EDQFS) [50] measures how uniformly information is distributed in the quantum feature space. It is based on the PCA eigenvalue spectrum and captures the intrinsic dimensionality of the embedded data. A high EDQFS suggests a flat eigenvalue distribution and a more balanced use of the available Hilbert space dimensions.

The effective dimension is calculated as

$$d_{\text{Eff}} = \frac{(\sum_i \lambda_i)^2}{\sum_i \lambda_i^2}, \quad (7)$$

where λ_i are the PCA eigenvalues of the quantum-encoded dataset that was encoded by the feature map one used. The summation runs over all principal components, i.e., $i = 1, \dots, r$, where $r = \min(n, d)$ is the rank of the dataset with n samples and d features.

QMetric computes EDQFS by performing PCA on the quantum features and evaluating the above formula. This metric complements FMCR by indicating how efficiently the encoded dimensions are utilized, helping to diagnose over- or underspread feature distributions [51].

3.2.3. Quantum Layer Activation Diversity

Quantum Layer Activation Diversity (QLAD) [52] evaluates the diversity of measurement outcomes across samples in the quantum feature space. It is based on the variance of probability distributions obtained from quantum measurements, reflecting how varied the output activations are for different inputs.

The metric is defined as

$$\text{QLAD} = \frac{1}{n} \sum_{i=1}^n \text{Var}(p_i), \quad (8)$$

where p_i is the measurement probability distribution for the i -th sample and n is the number of samples.

In QMetric, QLAD is computed by estimating the variance across each sample's probability vector and averaging the results. Low QLAD may signal that the circuit is collapsing inputs into narrow output distributions, which could hinder expressivity and generalization capabilities of the model.

3.2.4. Quantum Output Sensitivity

Quantum Output Sensitivity (QOS) [53] measures how sensitive a quantum model's output is to small perturbations in the input. It captures robustness and smoothness of the mapping from classical data to quantum measurements. A low QOS implies a stable, noise-tolerant model, while a high value may indicate fragility or sharp decision boundaries.

The metric is calculated as

$$\text{QOS} = \mathbb{E} \left[\frac{\|f(x + \epsilon) - f(x)\|^2}{\|\epsilon\|^2} \right], \quad (9)$$

where $f(x)$ and $f(x+\epsilon)$ are the quantum model outputs for the original and perturbed inputs, respectively. Here, $\mathbb{E}$ denotes the empirical average over a batch of perturbation vectors ϵ , typically sampled from a zero-mean isotropic Gaussian distribution.

In QMetric, QOS is evaluated by generating perturbed versions of inputs, computing the model output differences, and normalizing by the squared perturbation norms. This metric is useful for analyzing encoding smoothness, adversarial stability, and overall model resilience.

Table 2

Summary of QMetric Metrics

Category	Metric	Purpose
Quantum Circuit	QCE, QCF	Expressibility and noise robustness of circuits.
	QLR	Balance of local vs. entangling gates.
	EEE, QMI	Entanglement and intra-circuit correlation.
Feature Space	FMCR, EDQFS	Data compression and feature spread.
	QLAD	Diversity of quantum activations.
	QOS	Sensitivity to small input changes.
Training Dynamics	TSI, TEI	Stability and efficiency of convergence.
	QGN, BPI	Gradient health and barren plateau risk.
	RQLSI, r-QTEI	Relative diagnostics vs classical models.

3.3. Training Dynamics

QMetric also tracks training behavior using the *Training Stability Index* (TSI), which compares variability in training and validation losses, and the *Training Efficiency Index* (TEI), which measures epochs needed to reach a target accuracy relative to model size.

Quantum-specific diagnostics include the *Quantum Gradient Norm* (QGN) and *Barren Plateau Indicator* (BPI), both of which expose vanishing gradients linked to deep or poorly initialized circuits. To compare hybrid and classical models, QMetric implements relative metrics, such as *RQLSI* and *r-QTEI*, to quantify differences in training efficiency and stability under aligned conditions. Together, these metrics can support targeted diagnosis of underperformance, guide ansatz design, and enable meaningful evaluation across model types.

3.3.1. Training Stability Index

Training Stability Index (TSI) [54] quantifies the variability in training and validation losses near convergence. It measures how consistently the model performs in the final training phase by comparing the standard deviation of losses over the last 10% of epochs. This percentage will be up to the user in future versions of QMetric.

The metric is defined as

$$\text{TSI} = \frac{\sigma_{\text{val}}}{\sigma_{\text{train}}}, \quad (10)$$

where σ_{train} and σ_{val} denote the standard deviations of training and validation losses, respectively. Here, "losses" refer to the recorded values of the loss function (e.g., cross-entropy or MSE) over training and validation batches during training.

A low TSI indicates stable and consistent generalization, while a high value may reveal overfitting or noisy training dynamics. In QMetric, TSI is computed by evaluating standard deviations over the tail segment of the loss curves, i.e., we take into account the standard deviation of the last n outputs of the loss function.

3.3.2. Training Efficiency Index

Training Efficiency Index (TEI) [55] measures how quickly a model reaches a high level of performance relative to its size. It is defined as the ratio between the epoch at which validation accuracy first reaches a minimum $t = 90\%$ and the number of trainable parameters. While the value 90% is usually used, this threshold will be up to the user in future versions of QMetric.

Formally,

$$\text{TEI} = \frac{\text{epoch}_{\text{acc}_{\text{val}} \geq t}}{N_{\text{params}}}, \quad (11)$$

where $\text{epoch}_{\text{acc}_{\text{val}} \geq t}$ is the earliest epoch in which accuracy reaches desired threshold t and N_{params} is the total parameter count.

Lower TEI values indicate faster convergence per parameter, making this metric a useful tool for evaluating training efficiency across differently sized architectures.

3.3.3. Quantum Gradient Norm

Quantum Gradient Norm (QGN) [56] measures the magnitude of gradients associated with quantum circuit parameters. It reflects the overall strength of parameter updates and can signal the presence of vanishing or exploding gradients.

The metric is defined as

$$\text{QGN} = \left\| \nabla_{\theta_q} \mathcal{L} \right\|_2, \quad (12)$$

where θ_q denotes the quantum parameters and $\mathcal{L}$ is the training loss, i.e., the value of the loss function on the training dataset.

In QMetric, gradients are extracted from the backpropagation step and concatenated for L2 norm calculation

$$\|\nabla \mathcal{L}\|_2 = \left\| \left[\frac{\partial \mathcal{L}}{\partial \theta_1}, \frac{\partial \mathcal{L}}{\partial \theta_2}, \dots, \frac{\partial \mathcal{L}}{\partial \theta_n} \right] \right\|_2, \quad (13)$$

where $\mathcal{L}$ is the loss function and $\{\theta_i\}_{i=1}^n$ are the trainable parameters of the hybrid quantum-classical model.

Low QGN may indicate a barren plateau or excessively deep circuits [57].

3.3.4. Barren Plateau Indicator

Barren Plateau Indicator (BPI) [58] estimates whether a model suffers from barren plateaus by evaluating the average squared magnitude of quantum gradients. This captures the extent of vanishing gradients during optimization.

It is defined as

$$\text{BPI} = \mathbb{E} \left[\left\| \nabla_{\theta_q} \mathcal{L} \right\|^2 \right], \quad (14)$$

where $\mathbb{E}$ is the same as in eq. (9) and the rest of the parameters are as in eq. (12).

Values near zero suggest that gradients are vanishing, which can hinder effective training. In QMetric, BPI is computed over the flattened list of quantum gradients and averaged into a final value, making it an efficient early diagnostic tool during model tuning [59].

3.3.5. Relative Quantum Layer Stability Index

Relative Quantum Layer Stability Index (RQLSI) [60] compares the training stability of hybrid quantum-classical models to that of purely classical ones using the TSI metric. It helps quantify whether introducing quantum layers improves or worsens loss stability.

Formally,

$$\text{RQLSI} = \frac{\text{TSI}_{\text{hybrid}}}{\text{TSI}_{\text{classical}}}, \quad (15)$$

where $\text{TSI}_{\text{hybrid}}$ and $\text{TSI}_{\text{classical}}$ are the training stability indices for the hybrid and classical models, respectively.

A value less than 1 suggests that the quantum-enhanced model is more stable during training. This metric can support empirical comparison between model types under matched conditions.

3.3.6. Relative Quantum Training Efficiency Index

Relative Quantum Training Efficiency Index (r-QTEI) [61] evaluates whether a hybrid model trains more efficiently than a classical counterpart by comparing their respective TEI scores.

It is defined as

$$\text{r-QTEI} = \frac{\text{TEI}_{\text{hybrid}}}{\text{TEI}_{\text{classical}}}, \quad (16)$$

where $\text{TEI}_{\text{hybrid}}$ and $\text{TEI}_{\text{classical}}$ are the training efficiency indices.

A value below 1 means the hybrid model reaches target performance faster relative to its parameter size. This metric can support head-to-head benchmarking of model variants in practical scenarios.

4. Case Study: Hybrid vs Classical on MNIST

To illustrate the diagnostic capabilities of QMetric, we evaluate a hybrid quantum-classical neural network against a classical baseline using a binary classification task on the MNIST dataset [62]. This case study provides a practical scenario where quantum neural networks are tested under realistic constraints. We describe the model architectures, data pipeline, training configuration, and metric-driven analysis.

4.1. Hybrid Model and its Training Parameters

The hybrid model connects a parameterized quantum circuit with a classical output layer to perform binary classification. The quantum component is implemented using Qiskit's `EstimatorQNN` and is connected to PyTorch via the `TorchConnector`, allowing seamless integration with PyTorch's autograd system.

The quantum circuit is composed of a feature map and an ansatz. The feature map is a `ZZFeatureMap` with one repetition that encodes classical inputs into quantum states. The ansatz is a `RealAmplitudes` circuit that introduces trainable parameters and entanglement. The latter is repeated three times. `RealAmplitudes` circuits are composed into a single parameterized circuit, which is then used to define the quantum neural network. The quantum component outputs a single expectation value, which is passed through a trainable classical linear layer followed by a sigmoid activation. The full model maps input vector x to output $\sigma(W \cdot \text{QNN}(x) + b)$ where W and b are trainable classical parameters.

To match the number of qubits in the circuit, the MNIST images are projected into a lower-dimensional space using principal component analysis. The original 784-dimensional vectors are reduced to three components. This projection ensures compatibility with a three-qubit quantum circuit while preserving as much variance as possible.

The dataset is constructed by filtering the MNIST training set to include only samples corresponding to digits 0 and 1. From this filtered subset, the first 500 examples are selected to simulate a small-data regime. The images are flattened into vectors, normalized, and then transformed using [Principal Component Analysis \(PCA\)](#) to produce a dataset suitable for quantum encoding.

The hybrid model is trained for 30 epochs using the Adam optimizer with a learning rate of 0.01. Binary cross-entropy is used as the loss function. Training and validation losses are tracked at each epoch along with validation accuracy. Additionally, gradients concerning the quantum parameters are collected to enable computation of metrics such as the quantum gradient norm and barren plateau indicator.

After training, the quantum outputs are evaluated using `QMetric`. Metrics such as the feature map compression ratio, effective dimension of the quantum feature space, layer activation diversity, and output sensitivity are computed from the post-quantum activations. Quantum circuit diagnostics such as expressibility, locality ratio, entanglement entropy, mutual information, and noise robustness are also evaluated. This model provides a complete use case for applying `QMetric` during model selection, architectural tuning, and training analysis.

4.2. Classical Baseline

The classical baseline model is a fully connected neural network designed to match the input dimensionality and output behavior of the hybrid model. It takes as input the same three-dimensional data produced by PCA and outputs a binary classification probability using a sigmoid activation.

The architecture consists of an input layer with three neurons, a hidden layer with ten neurons using the ReLU activation function, and an output layer with one neuron followed by a sigmoid activation. The network approximates a function $f(x) = \sigma(W_2 \cdot \text{ReLU}(W_1 x + b_1) + b_2)$ where x is the PCA-reduced input vector and W_1 , W_2 , b_1 , and b_2 are trainable parameters.

The model is trained using the same subset of the MNIST dataset as the hybrid model. The inputs are 500 grayscale images corresponding to digits 0 and 1, flattened and reduced to three principal components. The preprocessing pipeline is identical, ensuring a fair comparison in terms of input dimensionality and task complexity.

The training procedure mirrors that of the hybrid model. The optimizer is Adam with a learning rate of 0.01, the loss function is binary cross-entropy, and the number of training epochs is set to 30. At each epoch, training loss, validation loss, and validation accuracy are recorded to allow for direct comparison of convergence dynamics, generalization performance, and learning stability.

This classical model serves as a baseline for interpreting the added value or limitations of quantum components under identical data, dimensionality, and optimization conditions. It enables a controlled




Figure 1: Hybrid quantum-classical model architecture. Classical inputs are encoded into quantum states via a feature map and processed by a parameterized circuit. The classical output from the quantum component is passed to a classical linear layer and sigmoid activation for binary classification.

analysis of the effects of quantum layers on expressivity, robustness, and trainability using QMetric’s evaluation framework.




Figure 2: Architecture of the baseline classical neural network. A fully connected feedforward model processes PCA-reduced MNIST inputs to perform binary classification between digits 0 and 1.

4.3. Quantum Circuit Metrics

Table 3 summarizes the metrics that characterize the structure, expressibility, and robustness of the quantum circuit used in the hybrid model.

The quantum circuit demonstrates high expressibility ($QCE = 0.939$), suggesting that it explores a diverse set of quantum states across the Hilbert space. The fidelity score ($QCF = 1.000$) confirms

Table 3
Quantum Circuit Metrics

Metric	Value	Interpretation
Quantum Circuit Expressibility (QCE)	0.939	High expressibility, close to Haar-random.
Quantum Volume Contribution (QVC)	0.1905	Moderate depth, penalizing expressive capacity.
Quantum Circuit Fidelity (QCF)	1.000	Perfect fidelity under simulation noise.
Quantum Locality Ratio (QLR)	0.6364	Balanced mix of local and entangling gates.
Effective Entanglement Entropy (EEE)	0.8345	High entanglement across qubit partitions.
Quantum Mutual Information (QMI)	1.6691	Strong total correlations between subsystems.

robustness to noise under simulation with a basic noise model. The locality ratio of 0.6364 indicates a well-balanced design between local and entangling operations. Entanglement is both substantial and well-structured, as shown by high values of $EEE = 0.8345$ and $QMI = 1.6691$, supporting rich correlations necessary for quantum information processing.

4.4. Feature Space Metrics

The geometry and structure of the quantum feature space are assessed through the metrics in Table 4, which evaluate compression, variance distribution, activation diversity, and sensitivity to perturbations.

Table 4
Quantum Feature Space Metrics

Metric	Value	Interpretation
Feature Map Compression Ratio (FMCР)	3.000	Strong compression from 3D input to 1D effective space.
Effective Dimension (EDQFS)	1.000	Variance concentrated in a single direction.
Quantum Layer Activation Diversity (QLAD)	0.000	Collapsed outputs, no diversity in activation patterns.
Quantum Output Sensitivity (QOS)	9.644	Highly sensitive to small input perturbations.

The feature map achieves perfect compression ($FMCР = 3.0$), indicating that all input variance is concentrated in one effective principal component. However, the effective dimension ($EDQFS = 1.0$) confirms that the quantum feature space lacks spread. Activation diversity is absent ($QLAD = 0.000$), signaling possible circuit over-regularization or symmetry that collapses measurement outputs. Meanwhile, the high sensitivity ($QOS = 9.644$) indicates the model reacts sharply to small perturbations, suggesting brittle or sharp decision boundaries.

4.5. Training Dynamics

Training dynamics of the hybrid and classical models are evaluated in Table 5. These metrics reflect convergence behavior, parameter efficiency, gradient stability, and vanishing gradient issues.

Table 5
Training Dynamics Metrics

Metric	Hybrid	Classical	Interpretation
Training Stability Index (TSI)	0.0025	0.0144	Hybrid is more stable near convergence.
Training Efficiency Index (TEI)	×	0.0000	Hybrid never reached 90% accuracy.
Quantum Gradient Norm (QGN)	0.458	—	Moderate gradient magnitude in last epoch.
Barren Plateau Indicator (BPI)	0.0175	—	Small but non-vanishing gradients.
Relative Stability (RQLSI)	0.1772	—	The Hybrid model shows lower variance.
Relative Efficiency (r-QTEI)	∞	—	Classical model is significantly faster to train.

The hybrid model exhibits lower validation loss variability in late training ($TSI = 0.0025$), indicating

consistent behavior, whereas the classical model converges quickly but shows slightly more fluctuation (TSI = 0.0144)

5. Outlook

QMetric provides a structured approach for diagnosing hybrid quantum-classical models beyond conventional performance metrics. It highlights key aspects such as training behavior, encoding robustness, and circuit design quality. Future developments will include migration to Qiskit's Estimator v2⁸, support for additional platforms like PennyLane, and expanded metric coverage for multi-class tasks and generative models.

6. Availability

All source code, examples, metric definitions, and plotting utilities are available at <https://gitlab.com/illesova.silvie.scholar/qmetric>. The repository includes a Conda environment file to reproduce the case study.

7. Acknowledgments

MB was supported by Italian Government (Ministero dell'Università e della Ricerca, PRIN 2022 PNRR)—cod. P2022SELA7: “RECHARGE: monitoRing, tEsting, and CHaracterization of performAnce Regressions”—D.D. n. 1205 del 28/7/2023. TR gratefully acknowledges the funding support by program “Excellence initiative—research university” for the AGH University in Kraków, as well as the ARTIQ project: UMO-2021/01/2/ST6/00004 and ARTIQ/0004/2021.

Declaration on Generative AI

During the preparation of this work, the authors used OpenAI ChatGPT (GPT-4) in order to: assist with text formatting, clarify structure, and edit grammar. After using these tools, the author reviewed and edited the content as needed and takes full responsibility for the publication's content.

References

- [1] J. Liu, K. H. Lim, K. L. Wood, W. Huang, C. Guo, H.-L. Huang, Hybrid quantum-classical convolutional neural networks, *Science China Physics, Mechanics & Astronomy* 64 (2021) 290311.
- [2] H.-Y. Chen, Y.-J. Chang, S.-W. Liao, C.-R. Chang, Deep q-learning with hybrid quantum neural network on solving maze problems, *Quantum Machine Intelligence* 6 (2024) 2.
- [3] Y. Zeng, H. Wang, J. He, Q. Huang, S. Chang, A multi-classification hybrid quantum neural network using an all-qubit multi-observable measurement strategy, *Entropy* 24 (2022) 394.
- [4] J. Biamonte, P. Wittek, N. Pancotti, P. Rebentrost, N. Wiebe, S. Lloyd, *Quantum machine learning*, *Nature* 549 (2017) 195–202.
- [5] M. Schuld, I. Sinayskiy, F. Petruccione, An introduction to quantum machine learning, *Contemporary Physics* 56 (2015) 172–185.
- [6] A. Senokosov, A. Sedykh, A. Sagingalieva, B. Kyriacou, A. Melnikov, *Quantum machine learning for image classification*, *Machine Learning: Science and Technology* 5 (2024) 015040.
- [7] W. Lu, Y. Lu, J. Li, A. Sigov, L. Ratkin, L. A. Ivanov, *Quantum machine learning: Classifications, challenges, and solutions*, *Journal of Industrial Information Integration* (2024) 100736.

⁸<https://quantum.cloud.ibm.com/docs/en/api/qiskit-ibm-runtime/estimator-v2>

- [8] C. A. Riofrio, O. Mitevski, C. Jones, F. Krellner, A. Vuckovic, J. Doetsch, J. Klepsch, T. Ehmer, A. Luckow, A characterization of quantum generative models, *ACM Transactions on Quantum Computing* 5 (2024) 1–34.
- [9] B. Coyle, M. Henderson, J. C. J. Le, N. Kumar, M. Paini, E. Kashefi, Quantum versus classical generative modelling in finance, *Quantum Science and Technology* 6 (2021) 024013.
- [10] R. Divya, J. D. Peter, Quantum machine learning: A comprehensive review on optimization of machine learning algorithms, in: 2021 Fourth International Conference on Microelectronics, Signals & Systems (ICMSS), IEEE, 2021, pp. 1–6.
- [11] N. A. AL Ajmi, M. Shoaib, Optimization strategies in quantum machine learning: A performance analysis, *Applied Sciences* 15 (2025) 4493.
- [12] V. Novák, I. Zelinka, V. Snášel, Optimization strategies for variational quantum algorithms in noisy landscapes, *arXiv preprint arXiv:2506.01715* (2025).
- [13] A. Bílek, J. Hlisnikovský, T. Bezděk, R. Kukulski, P. Lewandowska, Experimental study of multiple-shot unitary channels discrimination using the ibm q computers, *arXiv preprint arXiv:2505.17731* (2025).
- [14] P. Lewandowska, M. Beseda, Benchmarking gate-based quantum devices via certification of qubit von neumann measurements, *arXiv preprint arXiv:2506.03514* (2025).
- [15] V. Novák, I. Zelinka, L. Příbylová, L. Martínek, V. Benčurik, Predicting post-surgical complications with quantum neural networks: A clinical study on anastomotic leak, *arXiv preprint arXiv:2506.01708* (2025).
- [16] L. Wei, H. Liu, J. Xu, L. Shi, Z. Shan, B. Zhao, Y. Gao, Quantum machine learning in medical image analysis: A survey, *Neurocomputing* 525 (2023) 42–53.
- [17] S. Rani, P. K. Pareek, J. Kaur, M. Chauhan, P. Bhambri, Quantum machine learning in healthcare: Developments and challenges, in: 2023 IEEE International Conference on Integrated Circuits and Communication Systems (ICICACS), IEEE, 2023, pp. 1–7.
- [18] J. Thiayagalingam, M. Shankar, G. Fox, T. Hey, Scientific machine learning benchmarks, *Nature Reviews Physics* 4 (2022) 413–420.
- [19] M. Sajjan, J. Li, R. Selvarajan, S. H. Sureshbabu, S. S. Kale, R. Gupta, V. Singh, S. Kais, Quantum machine learning for chemistry and physics, *Chemical Society Reviews* 51 (2022) 6475–6573.
- [20] B. Huang, N. O. Symonds, O. A. von Lilienfeld, Quantum machine learning in chemistry and materials, *Handbook of Materials Modeling: Methods: Theory and Modeling* (2020) 1883–1909.
- [21] N. Bauer, K. Yeter-Aydeniz, G. Siopsis, Efficient quantum chemistry calculations on noisy quantum hardware, *arXiv preprint arXiv:2503.02778* (2025).
- [22] D. A. Fedorov, B. Peng, N. Govind, Y. Alexeev, Vqe method: a short survey and recent developments, *Materials Theory* 6 (2022) 2.
- [23] S. Illéssová, V. Novák, T. Bezděk, M. Beseda, C. Possel, Numerical optimization strategies for the variational hamiltonian ansatz in noisy quantum environments, *arXiv preprint arXiv:2505.22398* (2025).
- [24] C. Ciaramelletti, M. Beseda, M. Consiglio, L. Lepori, T. J. G. Apollaro, S. Paganelli, Detecting quasidegenerate ground states in topological models via the variational quantum eigensolver, *Phys. Rev. A* 111 (2025) 022437. URL: <https://link.aps.org/doi/10.1103/PhysRevA.111.022437>. doi:10.1103/PhysRevA.111.022437.
- [25] C. Zhang, L. Jiang, F. Chen, Qracle: A graph-neural-network-based parameter initializer for variational quantum eigensolvers, *arXiv preprint arXiv:2505.01236* (2025).
- [26] M. Beseda, S. Illéssová, S. Yalouz, B. Senjean, State-averaged orbital-optimized vqe: A quantum algorithm for the democratic description of ground and excited electronic states, *Journal of Open Source Software* 9 (2024) 6036.
- [27] S. Illéssová, M. Beseda, S. Yalouz, B. Lasorne, B. Senjean, Transformation-free generation of a quasi-diabatic representation from the state-average orbital-optimized variational quantum eigensolver, *Journal of Chemical Theory and Computation* (2025).
- [28] H. L. Tang, V. Shkolnikov, G. S. Barron, H. R. Grimsley, N. J. Mayhall, E. Barnes, S. E. Economou, qubit-adapt-vqe: An adaptive algorithm for constructing hardware-efficient ansätze on a quantum

processor, *PRX Quantum* 2 (2021) 020310.

- [29] K. M. Nakanishi, K. Mitarai, K. Fujii, Subspace-search variational quantum eigensolver for excited states, *Physical Review Research* 1 (2019) 033062.
- [30] M. K. Gupta, M. Beseda, P. Gawron, How quantum computing-friendly multispectral data can be?, in: *IGARSS 2022-2022 IEEE International Geoscience and Remote Sensing Symposium*, IEEE, 2022, pp. 4153–4156.
- [31] Y. Suzuki, H. Yano, Q. Gao, S. Uno, T. Tanaka, M. Akiyama, N. Yamamoto, Analysis and synthesis of feature map for kernel-based quantum classifier, *Quantum Machine Intelligence* 2 (2020) 9.
- [32] H. Kwon, H. Lee, J. Bae, Feature map for quantum data in classification, in: *2024 International Conference on Quantum Communications, Networking, and Computing (Q CNC)*, IEEE, 2024, pp. 41–48.
- [33] S. Gupta, R. Zia, Quantum neural networks, *Journal of Computer and System Sciences* 63 (2001) 355–383.
- [34] P. Rebentrost, M. Mohseni, S. Lloyd, Quantum support vector machine for big data classification, *Physical review letters* 113 (2014) 130503.
- [35] C. Wilson, J. Otterbach, N. Tezak, R. Smith, A. Polloreno, P. J. Karalekas, S. Heidel, M. S. Alam, G. Crooks, M. da Silva, Quantum kitchen sinks: An algorithm for machine learning on near-term quantum computers, *arXiv preprint arXiv:1806.08321* (2018).
- [36] V. Havlíček, A. D. Córcoles, K. Temme, A. W. Harrow, A. Kandala, J. M. Chow, J. M. Gambetta, Supervised learning with quantum-enhanced feature spaces, *Nature* 567 (2019) 209–212.
- [37] J.-G. Liu, L. Wang, Differentiable learning of quantum circuit born machines, *Physical Review A* 98 (2018) 062324.
- [38] D. F. Locher, L. Cardarelli, M. Müller, Quantum error correction with quantum autoencoders, *Quantum* 7 (2023) 942.
- [39] F. Zhang, J. Li, Z. He, H. Situ, Transformer for parameterized quantum circuits expressibility prediction, *arXiv e-prints* (2024) arXiv–2405.
- [40] D. Ter Haar, Foundations of statistical mechanics, *Reviews of Modern Physics* 27 (1955) 289.
- [41] R.-Q. Yang, Gravity duals of quantum distances, *Journal of High Energy Physics* 2021 (2021) 1–12.
- [42] S. Kullback, Kullback-leibler divergence, *Tech. Rep.* (1951).
- [43] P. E. Mendonça, R. d. J. Napolitano, M. A. Marchioli, C. J. Foster, Y.-C. Liang, Alternative fidelity measure between quantum states, *Physical Review A—Atomic, Molecular, and Optical Physics* 78 (2008) 052330.
- [44] A. Vadali, R. Kshirsagar, P. Shyamsundar, G. N. Perdue, Quantum circuit fidelity estimation using machine learning, *Quantum Machine Intelligence* 6 (2024) 1.
- [45] B. Bertini, P. Kos, T. Prosen, Operator entanglement in local quantum circuits i: Chaotic dual-unitary circuits, *SciPost Physics* 8 (2020) 067.
- [46] P. Boes, J. Eisert, R. Gallego, M. P. Müller, H. Wilming, Von neumann entropy from unitarity, *Physical review letters* 122 (2019) 210402.
- [47] B. Schumacher, M. D. Westmoreland, Quantum mutual information and the one-time pad, *Physical Review A—Atomic, Molecular, and Optical Physics* 74 (2006) 042305.
- [48] G. De Tomasi, S. Bera, J. H. Bardarson, F. Pollmann, Quantum mutual information as a probe for many-body localization, *Physical Review Letters* 118 (2017) 016804.
- [49] A. Maćkiewicz, W. Ratajczak, Principal components analysis (pca), *Computers & Geosciences* 19 (1993) 303–342.
- [50] J. Crossley, A. Nerode, Effective dimension, *Journal of Algebra* 41 (1976) 398–412.
- [51] M. Schuld, N. Killoran, Quantum machine learning in feature hilbert spaces, *Physical review letters* 122 (2019) 040504.
- [52] B. Xie, Y. Liang, L. Song, Diverse neural network learns true target functions, in: *Artificial Intelligence and Statistics*, PMLR, 2017, pp. 1216–1224.
- [53] D. Liu, L. Y. Wu, B. Li, F. Boussaid, M. Bennamoun, X. Xie, C. Liang, Jacobian norm with selective input gradient regularization for interpretable adversarial defense, *Pattern Recognition* 145 (2024) 109902.

- [54] N. Chandramoorthy, A. Loukas, K. Gatmiry, S. Jegelka, On the generalization of learning algorithms that do not converge, *Advances in Neural Information Processing Systems* 35 (2022) 34241–34257.
- [55] R. Livni, S. Shalev-Shwartz, O. Shamir, On the computational efficiency of training neural networks, *Advances in neural information processing systems* 27 (2014).
- [56] N. Galegale, C. I. Shimabukuro, et al., Deep learning applied to stock prices: Epoch adjustment in training an lstm neural network, *International Journal of Business and Management* 19 (2024).
- [57] A. Gilyén, S. Arunachalam, N. Wiebe, Optimizing quantum optimization algorithms via faster quantum gradient computation, in: *Proceedings of the Thirtieth Annual ACM-SIAM Symposium on Discrete Algorithms*, SIAM, 2019, pp. 1425–1444.
- [58] M. Larocca, P. Czarnik, K. Sharma, G. Muraleedharan, P. J. Coles, M. Cerezo, Diagnosing barren plateaus with tools from quantum optimal control, *Quantum* 6 (2022) 824.
- [59] J. R. McClean, S. Boixo, V. N. Smelyanskiy, R. Babbush, H. Neven, Barren plateaus in quantum neural network training landscapes, *Nature communications* 9 (2018) 4812.
- [60] Y. Liu, Q. Li, J. Tan, Y. Shi, L. Shen, X. Cao, Understanding the stability-based generalization of personalized federated learning, in: *The Thirteenth International Conference on Learning Representations*, ????
- [61] M. Tan, Q. Le, Efficientnetv2: Smaller models and faster training. arxiv 2021, *arXiv preprint arXiv:2104.00298* 5 (2021).
- [62] L. Deng, The mnist database of handwritten digit images for machine learning research, *IEEE Signal Processing Magazine* 29 (2012) 141–142.