

Photovoltaic Power Forecasting Using Quantum Machine Learning

Machine learning

Photovoltaic power forecasting using quantum machine learning

Asel Sagingalieva,¹ Stefan Komorniyk,² Arsenii Senokosov,¹ Ayush Joshi,¹ Alexander Sedykh,¹ Christopher Mansell,¹ Olga Tsurkan,¹ Karan Pinto,¹ Markus Pflictsch,¹ and Alexey Melnikov¹

¹*Terra Quantum AG, 9000 St. Gallen, Switzerland*

²*HAKOM Time Series GmbH, 1230 Vienna, Austria*

Predicting solar panel power output is crucial for advancing the energy transition but is complicated by the variable and non-linear nature of solar energy. This is influenced by numerous meteorological factors, geographical positioning, and photovoltaic cell properties, posing significant challenges to forecasting accuracy and grid stability. Our study introduces a suite of solutions centered around hybrid quantum neural networks designed to tackle these complexities. The first proposed model, the Hybrid Quantum Long Short-Term Memory, surpasses all tested models by over 40% lower mean absolute and mean squared errors. The second proposed model, Hybrid Quantum Sequence-to-Sequence neural network, once trained, predicts photovoltaic power with 16% lower mean absolute error for arbitrary time intervals without the need for prior meteorological data, highlighting its versatility. Moreover, our hybrid models perform better even when trained on limited datasets, underlining their potential utility in data-scarce scenarios. These findings represent a stride towards resolving time series prediction challenges in energy power forecasting through hybrid quantum models, showcasing the transformative potential of quantum machine learning in catalyzing the renewable energy transition.

I. INTRODUCTION

Electricity generation prediction, especially for photovoltaic (PV) systems, is a crucial tool for renewable energy adoption [1, 2]. The global economy must radically reduce emissions to stay within the 1.5°C pathway (Paris Agreement) and the transition to renewable energy sources is necessary to achieve these objectives [3]. According to the IEA, solar PV’s installed power capacity is poised to surpass that of coal by 2027, becoming the largest in the world.

Accurate PV power forecasts are vital for multiple facets of the energy industry such as long-term investment planning, regulatory compliance for avoiding penalties, and renewable energy management across storage, transmission, and distribution activities. Several studies show that an increase in forecasting accuracy reduces electricity generation from conventional sources. Increased accuracy also reduces operating costs of systems through reducing the uncertainty of PV power generation [4]. They support improving the stability and sustainability of the power grid through optimizing power flow and counteracting solar power’s intermittent nature [5]. Such predictions are foundational in increasing the economic viability and improving the adoption of solar energy as they inform pricing and economic dispatch strategies, bolster competitiveness and over time reduce reliance on reserve power. Additionally, they assist in managing energy storage effectively and integrating PV systems into the power grid [6], which is essential for the enduring success of renewable energy solutions [7].

Traditional methods for predicting PV power have primarily relied on statistical models, machine learning algorithms, or a blend of both [8]. These approaches encompass a diverse toolkit, ranging from time series forecasting and artificial neural networks [1, 9, 10], to support vector machines [11, 12], k-nearest neighbor meth-

ods [13], and random forest models [14]. However, the intermittent and non-linear nature of solar power generation, influenced by a wide range of meteorological factors, poses a significant challenge to the performance of these conventional models [15].

In light of these challenges, quantum machine learning (QML) emerges as a promising avenue. This rapidly evolving field, which melds the principles of quantum mechanics with classical machine learning [16–19], can offer enhanced capabilities for improving the forecasting accuracy of time series tasks [20], including PV power generation [21]. QML’s potential arises from quantum features like superposition and entanglement, promising exponential speedups in certain tasks [22]. Moreover, QML algorithms produce inherently probabilistic results, aptly suited for prediction tasks, and they may potentially function within an exponentially larger search space, amplifying their efficacy [23–26]. Nonetheless, implementing quantum algorithms bears its own set of challenges, such as the need for error correction and sensitivity to external interference [27]. Yet, in spite of these challenges, hybrid quantum-classical models, especially hybrid quantum neural networks (HQNNs), have showcased their potential in diverse industrial realms, including healthcare [28–30], energy [21, 31], aerospace [32], logistics [33] and automotive [34] industries.

In this article, we present three types of hybrid quantum models as potential solutions for PV power forecasting. We assess the performance of our proposed models using a publicly accessible dataset, encompassing a comprehensive array of meteorological variables as well as hourly mean PV power measurements spanning a 21-month period. This dataset, along with the data preprocessing and analytical methodologies employed, is described in detail in Section II A.

Our first proposed HQNN architecture, articulated in Section II B, incorporates classical fully connected lay-

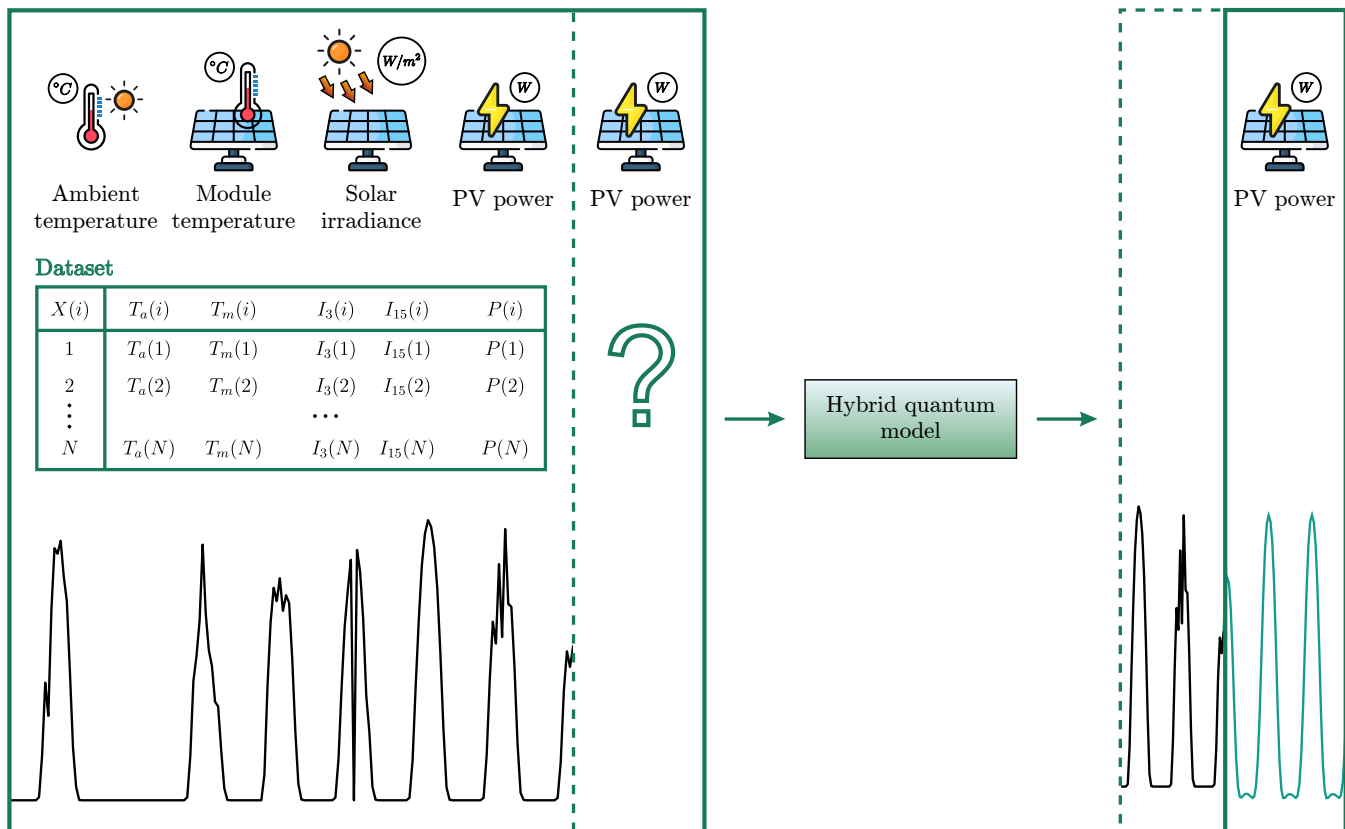


FIG. 1: The input to a hybrid quantum model is presented as a chronological data table, documenting hourly meteorological parameters including ambient temperature (T_a), module temperature (T_m), and solar irradiance (I_3 , I_{15}), alongside the mean PV power output (P). The model is designed to leverage this data to generate predictions of PV power output for a short-term forecast aimed at near-future output, typically the next hour, and a long-term forecast that extends to a broader temporal horizon.

ers with a vanilla variational repetitive quantum layer (VVRQ). Our second model, delineated in Section II C, constitutes a hybrid quantum adaptation of the classical recurrent neural network, termed the Hybrid Quantum Long Short-Term Memory with quantum depth-infused layer (HQLSTM). While the first two models can predict the power for a certain hour ahead, the third model, presented in Section II D, a Hybrid Quantum Sequence-to-Sequence Neural Network with quantum depth-infused layer, HQSeq2Seq, after training is capable of forecasting PV power for arbitrary time intervals without requiring prior meteorological data.

Remarkably, despite having fewer parameters, our hybrid quantum models outperform their classical counterparts in terms of more accurate predictions, including trained on a reduced dataset. We summarize our conclusions and outline future research directions in Section III.

II. RESULTS

The application of HQNNs in addressing time series prediction challenges, specifically in forecasting PV

power output offers several advantages. Primarily, their capability to operate within an exponentially larger computational search space enables them to efficiently capture intricate data patterns and relationships [35]. This feature not only enhances forecast accuracy [17, 36] but also streamlines the learning process, requiring fewer iterations for model optimization [37]. Furthermore, the inherent capacity of quantum technologies to manage the uncertainty and noise ubiquitous in data offers more resilient and trustworthy predictions [22]. This is particularly pertinent to power forecasting, given the inherent noise in meteorological data. Recent research also suggests that quantum models can be represented as partial Fourier series, positioning them as potential universal function approximators [38], thereby broadening their applicability and efficacy in predictive tasks.

In terms of architecture, an HQNN is an amalgamation of classical and quantum components. The classical segments may consist of fully connected layers, convolutional layers, or recurrent layers, while the quantum segments are typically represented by variational quantum circuits (VQCs) or their contemporary modifications [39, 40].

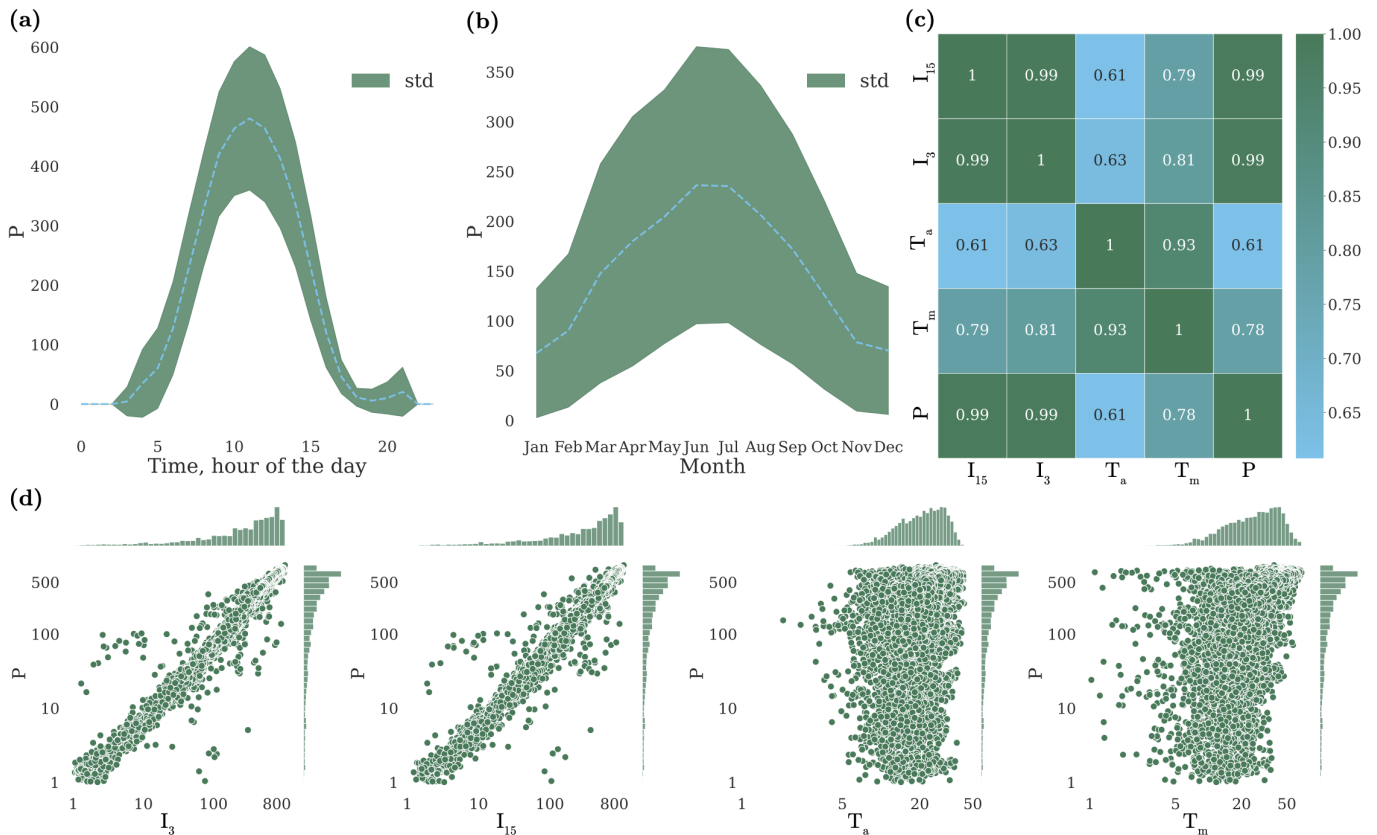


FIG. 2: (a) Mean and standard deviation of the PV power value for each hour of the day. (b) Mean and standard deviation of the PV power value for each month of the year. The plot shows the PV power reaching maximal values in June and July. (c) Correlation matrix of input features. (d) Joint distribution of features.

A. Dataset

To underscore the advantages of hybrid quantum models using empirical evidence, we selected a publicly accessible dataset [41] from a conventional generation plant situated in the Mediterranean region. This dataset not only provides comprehensive data but also allows benchmarking with results from various algorithms available in literature. A comparative analysis of our model’s predictions and those from the study by [10] is provided in Section II E. The dataset, presented as a numerical table showcased in Fig. 1, encompasses variables like hourly mean ambient temperature (T_a), hourly mean module temperature (T_m), hourly mean solar irradiance recorded on two tilted planes with tilt angles of 3 and 15 degrees (I_3 , I_{15}), and hourly mean PV power (P) spanning 21 months, accounting for more than 500 days.

Beyond the scope of constructing models for predicting the output of PV panels, this dataset’s utility extends to other applications. It aids in planning distributed battery energy storage systems [42], devising novel energy collection systems [43], and researching the degradation patterns of photovoltaic panels [44]. The dataset’s multifaceted applicability emphasizes its significance.

To ensure the validity and precision of the data, meticulous preprocessing and analysis were undertaken. We discovered approximately 20 anomalies in the original dataset. To maintain a continuous timeline, missing data points were replaced with the arithmetic mean of the preceding and succeeding day’s values. Additionally, data corresponding to the date “12/31/13” was excluded as it contained all-zero values, suggesting an error in data collection. As a result, we obtained an uninterrupted dataset ranging from 4:00 AM on “3/5/12” to 12:00 AM on “12/30/13”.

Additional in-depth analysis of the dataset was also conducted for a more nuanced understanding. Fig. 2(a) delineates the hourly distribution of PV power across the entire period. As expected, peak PV power values occur during midday, whereas night time values plummet to zero. Fig. 2(b) portrays monthly PV power fluctuations, which are more volatile compared to daily patterns, likely attributable to the limited number of full-year periods in the dataset. Fig. 2(c) presents a correlation matrix for the dataset features, identifying solar irradiances I_3 and I_{15} as the features most correlated with PV power. Finally, the joint distribution of dataset features depicted in Fig. 2(d) further confirms that solar intensity is the

feature most highly correlated with PV power.

B. HQNN

This section introduces our first proposed model, referred to as the HQNN. As illustrated in Fig. 1, the model accepts weather data spanning 24 consecutive hours as its input. The output is a prediction of the PV power for the upcoming 25th hour. The HQNN presented at the Fig. 3(a) is a combination of classical fully-connected layers, in our case with 120, 17 and 8 neurons, and a VVRQ layer, which is a VQC, consisting of q qubits and d repetitions of variational layers, each distinguished by unique weights. The choice of 120 neurons is methodical: the model ingests 5 distinct features for each of the 24 hours, resulting in a total of $120 = 5 * 24$. The determination of the remaining parameters stemmed from an extensive hyperparameter optimization process, detailed in the subsequent sections.

Initially, every qubit in the VVRQ layer is set to the state $|0\rangle$. We subsequently encode the classical data by converting it into rotation angles around one of the X , Y , Z axes using R_x , R_y , R_z gates respectively. This conversion employs the angle embedding technique [45]. For each qubit, the rotation angle, denoted by x_j , is determined by the j -th component of the input vector.

Following this, the variational layer is applied, which can either utilize “basic” or “strong” entanglements. For the “basic” entanglement, each qubit undergoes a rotation by an angle w_j^i around the X axis, subsequently followed by a layer of CNOT gates [46]. Conversely, for the “strong” entanglement, each qubit is sequentially rotated by the angles $w_{ji}^{(Z_1)}$, $w_{ji}^{(Y_2)}$, and $w_{ji}^{(Z_3)}$ around the Z , Y , and Z axes, respectively. This sequence is then followed by a layer of CNOT gates. In both cases, the variables i and j play crucial roles in determining the operations. The variable i signifies the particular wire to which the operation is applied, and it takes values from the set $1, 2, \dots, q$. Meanwhile, the variable j represents the number of variational layers and ranges from $1, 2, \dots, d$.

Lastly, all qubits are measured in Pauli- Z basis, yielding the classical vector $v \in \mathbb{R}^q$. This output serves as input for a subsequent classical fully-connected layer. This layer processes information from q neurons into 1 neuron that predicts the power value.

The proposed HQNN model will be compared with its classical analog – a Multilayer Perceptron (MLP) that consists of 4 fully connected layers with 120, 32, 3, 3, and 1 neurons. The number of neurons in each layer was selected by a hyperparameter optimization procedure, detailed in the subsequent sections.

C. HQLSTM

This section presents a description of our second hybrid model – HQLSTM, which is a hybrid analog of the classical LSTM model [47], with which predictions will be compared in the following sections. LSTM architectures have garnered significant attention in the realm of time series forecasting, including in predicting PV power [48–50].

HQLSTM models have proven themselves well for solving problems from various fields. Examples of successful use of this model are the tasks of natural language processing [51], the detection of software vulnerabilities [52], and predicting solar radiation [53].

In this proposed model we added a quantum layer to each of the LSTM gates [54]. Let’s take a closer look at our implementation, depicted in Fig 3(b). The input to the model is:

1. The current step information, represented by a green circle, $x(t)$. This is a tensor of size 5, reflecting the five features for an hour, which include meteorological data and the PV power itself.
2. The information from the previous step, denoted by a purple circle, $h(t - 1)$. It consists of a tensor of size h_{dim} . For the initial step, this is simply a zero vector.

These inputs are processed through classical fully-connected layers to yield vectors with a uniform dimension of $4n_q$. These vectors are then concatenated through a bitwise addition operation.

Subsequently, this concatenated vector is partitioned into four distinct groups, for the four gates of the LSTM cell. As illustrated in Fig. 3(b-c), each group is directed to the input of its corresponding quantum layer, symbolized by the QDI square.

The outputs from QDI layers are transformed via classical fully-connected layers to standardize their dimensions to h_{dim} . Following this, activation functions together with appropriate for each of the 4 gates transformations, similar to the classical LSTM, are applied to the outputs originating from the quantum layers. This processing culminates in the derivation of the new cell state $C(t)$ and the hidden state $h(t)$ vectors.

The process operates in a cyclical manner. For each iteration, the vector from the current time step, $x(t)$, and the hidden vector from the previous step, $h(t - 1)$, serve as inputs to the HQLSTM. This iterative process is executed as many times as the input width; in our case input width equals 24. Subsequently, all the hidden vectors are concatenated to produce a single composite vector. This vector is then processed through a fully-connected layer consisting of a single neuron, which outputs a value that predicts the PV power.

In our first proposed architecture, the HQNN, the quantum layer functioned as a vanilla layer, where variational layers were sequentially placed after the encoding

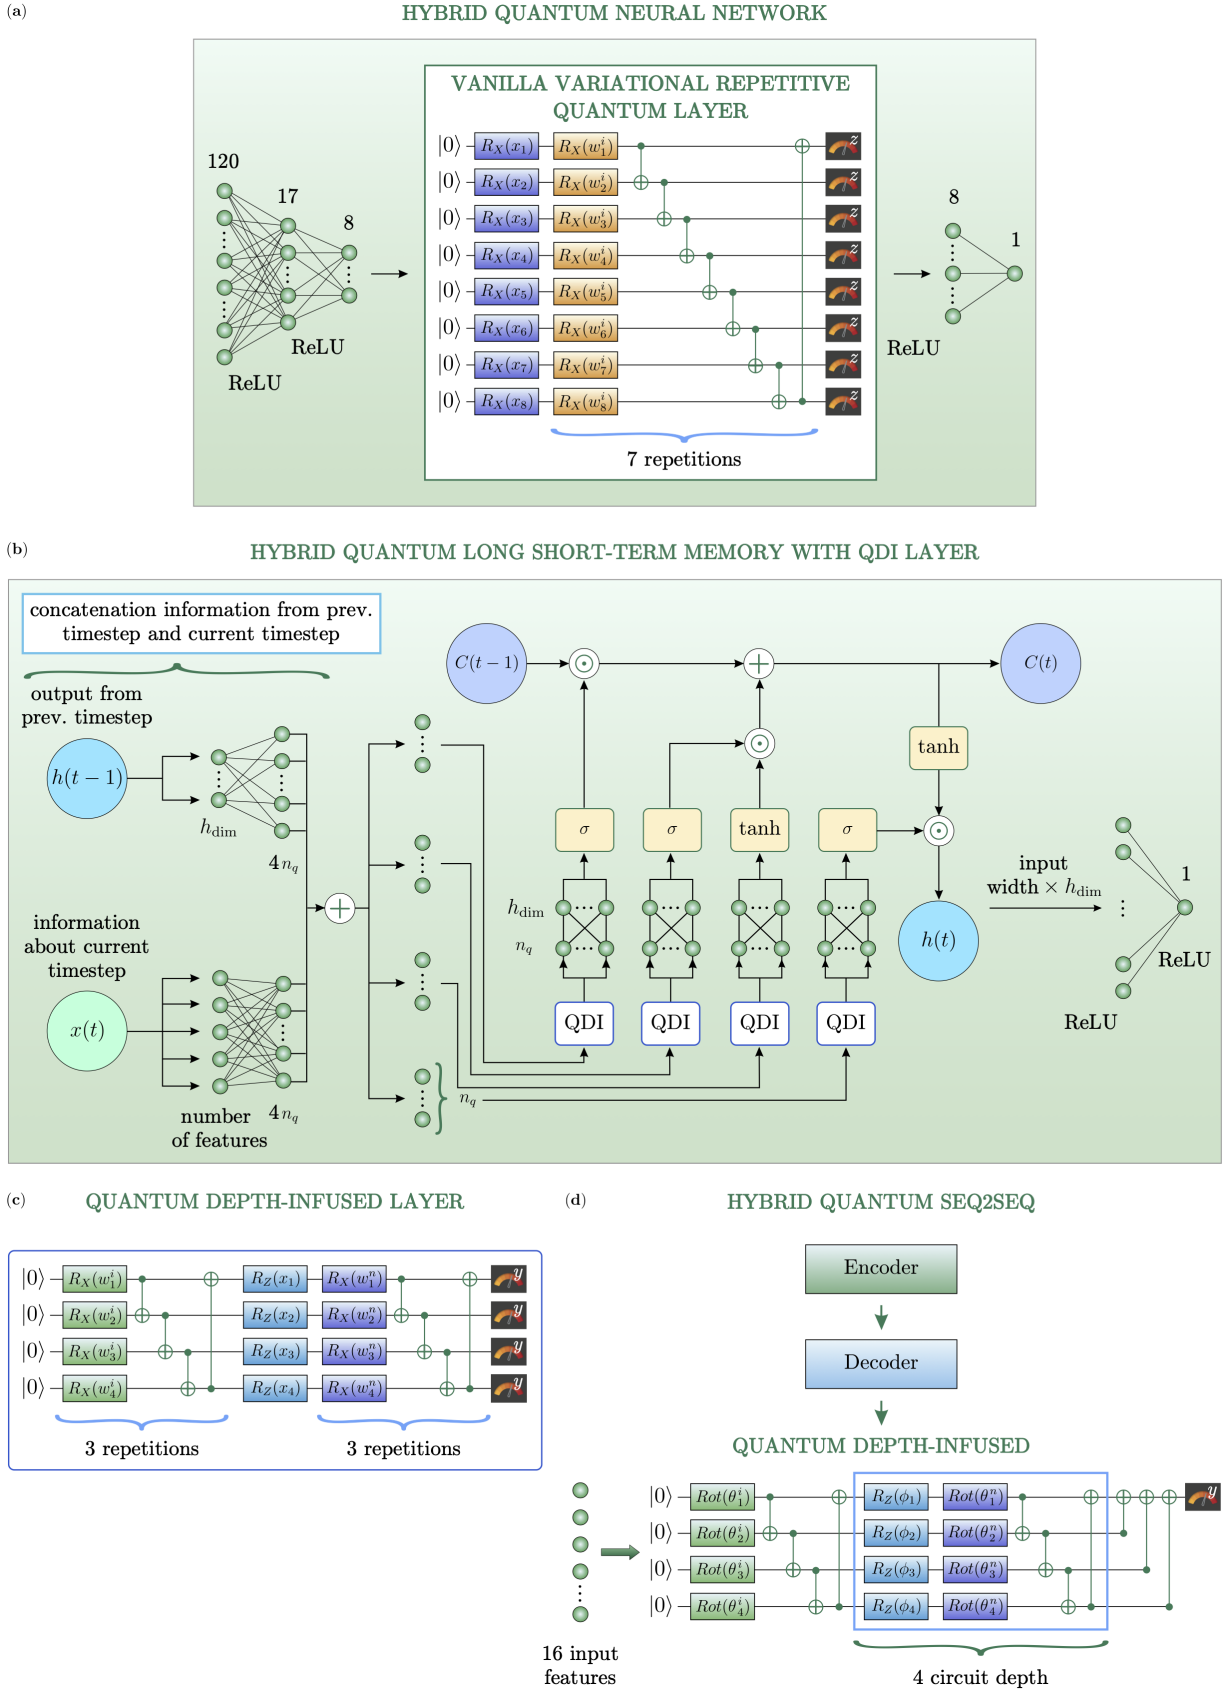


FIG. 3: The architectures of: (a) Hybrid Quantum Neural Network with VVRQ layer, (b) Hybrid Quantum Long Short-Term memory, (c) QDI layer used in HQLSTM model, (d) Hybrid Quantum Seq2Seq with QDI layer.

layer. In contrast, in the HQLSTM approach we used a QDI layer [29] as depicted in Fig. 3(b). Here, variational

layers are positioned multiple times before the encoding layer (green rectangular) and additionally (purple rectangular), within each encoding layer (blue rectangular) to increase the layer’s expressivity.

D. HQSeq2Seq

Here we present a hybrid version of the Sequence-to-Sequence (Seq2Seq) model, first introduced in [55]. Seq2Seq models are widely used in natural language processing tasks [56], where the length of input and output sequence is not pre-determined and can be variable. We can also apply the principle of Seq2Seq models to the power prediction task [57]. That means we can feed the neural network with time series with arbitrary length and prompt it to give us the forecast for any hours ahead. In this problem setting, the longer the input time series is, the better the model prediction is. The same applies to the required output length: the shorter it is, the easier it is for the model to generate the forecast.

The seq2Seq model is a type of encoder-decoder model. The encoder is given the entire input sequence, which it uses to generate a context vector. This vector is used as an input hidden state for the decoder, so it literally provides it with “context”, according to which the decoder will generate the forecast. Thereby, the hidden dimensions of the encoder and the decoder must match. The decoder creates the output sequence step by step. It starts with only one entry: the one which is the last known. Based on this entry and the context vector, the decoder generates the second entry and appends it to the existing one. Now, the obtained two-entry sequence is once again fed into the decoder to generate the third entry. Then, the cycle repeats until the length of the generated sequence matches the length requested by the user.

We create and compare two models with Seq2Seq architecture: the classical Seq2Seq and the hybrid model called HQSeq2Seq. Both of these models have identical LSTMs acting as encoders and decoders. In the classical model, the decoder’s hidden output vector is mapped to the “Power” value with a single linear layer, while in HQSeq2Seq it is processed by a QDI layer [29].

In the QDI layer, instead of attempting to use a qubit for each feature [37], we employed the data re-uploading technique [38, 58]. Specifically, we work with 4 qubits and structure them into a lattice of depth 4 (depicted as a blue big rectangular in Fig. 3(d)). Each of our 16 input features leading to the quantum layer is intricately encoded within this lattice. The first four features are mapped onto the initial depth, followed by the subsequent features in blocks of four. Encoding these classical features into the quantum domain, we adopt the “angle embedding” using R_z gate. This operation effectively translates the input vector into a quantum state that symbolizes the preceding classical layer’s data.

Entangling variational layers, signified by purple

squares, are interposed between every encoding layer, ensuring optimal Fourier accessibility. Each variational layer has two components: rotations governed by trainable parameters, and sequential CNOT gates. The rotations are implemented by quantum gates that metamorphose the encoded input in line with the variational parameters, while the CNOT operations handle the entanglement of the qubits, facilitating quantum superposition.

Each lattice depth, represented by each blue square, encompasses a variational layer (purple square). Moreover, prior to all encoding layers, we introduce a variational layer (designated by a green square) for enhanced model representation. Consequently, the total weight count in the quantum segment of our network is 20. In the measurement phase, except for the first qubit, all qubits execute a CNOT operation targeting the first qubit, ensuring the Y -measurement spans all qubits. Therefore, the quantum layer’s output serves as the power value prediction for a specific hour.

The input size of the encoder and decoder can differ, which is a substantial benefit. For instance, we can use all of the 5 features to create a context vector, but request to generate the forecast for only 1 feature. Exploiting this advantage, we will feed the Seq2Seq model with a window of all known features and demand the forecast only for the “Power” one.

For simplicity’s sake, we will train both models with fixed input and output length of 96 hours and then try to vary the length in the testing stage.

E. Training and results

In the study, six distinct models were employed for PV power prediction based on weather features: HQNN, MLP, HQLSTM, LSTM, HQSeq2Seq, and Seq2Seq. To train the models, the mean square error (MSE) was chosen as the loss function:

$$\text{MSE} = \frac{1}{N} \sum_{n=1}^N (x_n - y_n)^2,$$

where N is number of predictions, x denotes the predicted PV power, and y represents the actual PV power value.

To test the models, in addition to the MSE loss metric, we also used the mean absolute error (MAE), root mean squared error (RMSE) and variance account factor (VAF).

Here $\vec{x} = (x_1, x_2, \dots, x_N)$ and $\vec{y} = (y_1, y_2, \dots, y_N)$ represent vectors of predicted and target PV power values respectively, where N is the number of predicted values.

All the machine learning simulations for this study were conducted on CPUs, on the QMware cloud [59, 60] device. The classical part of our modeling was structured using the PyTorch library [61], while the quantum part was implemented using the PennyLane framework. Notably, PennyLane provides an assortment of

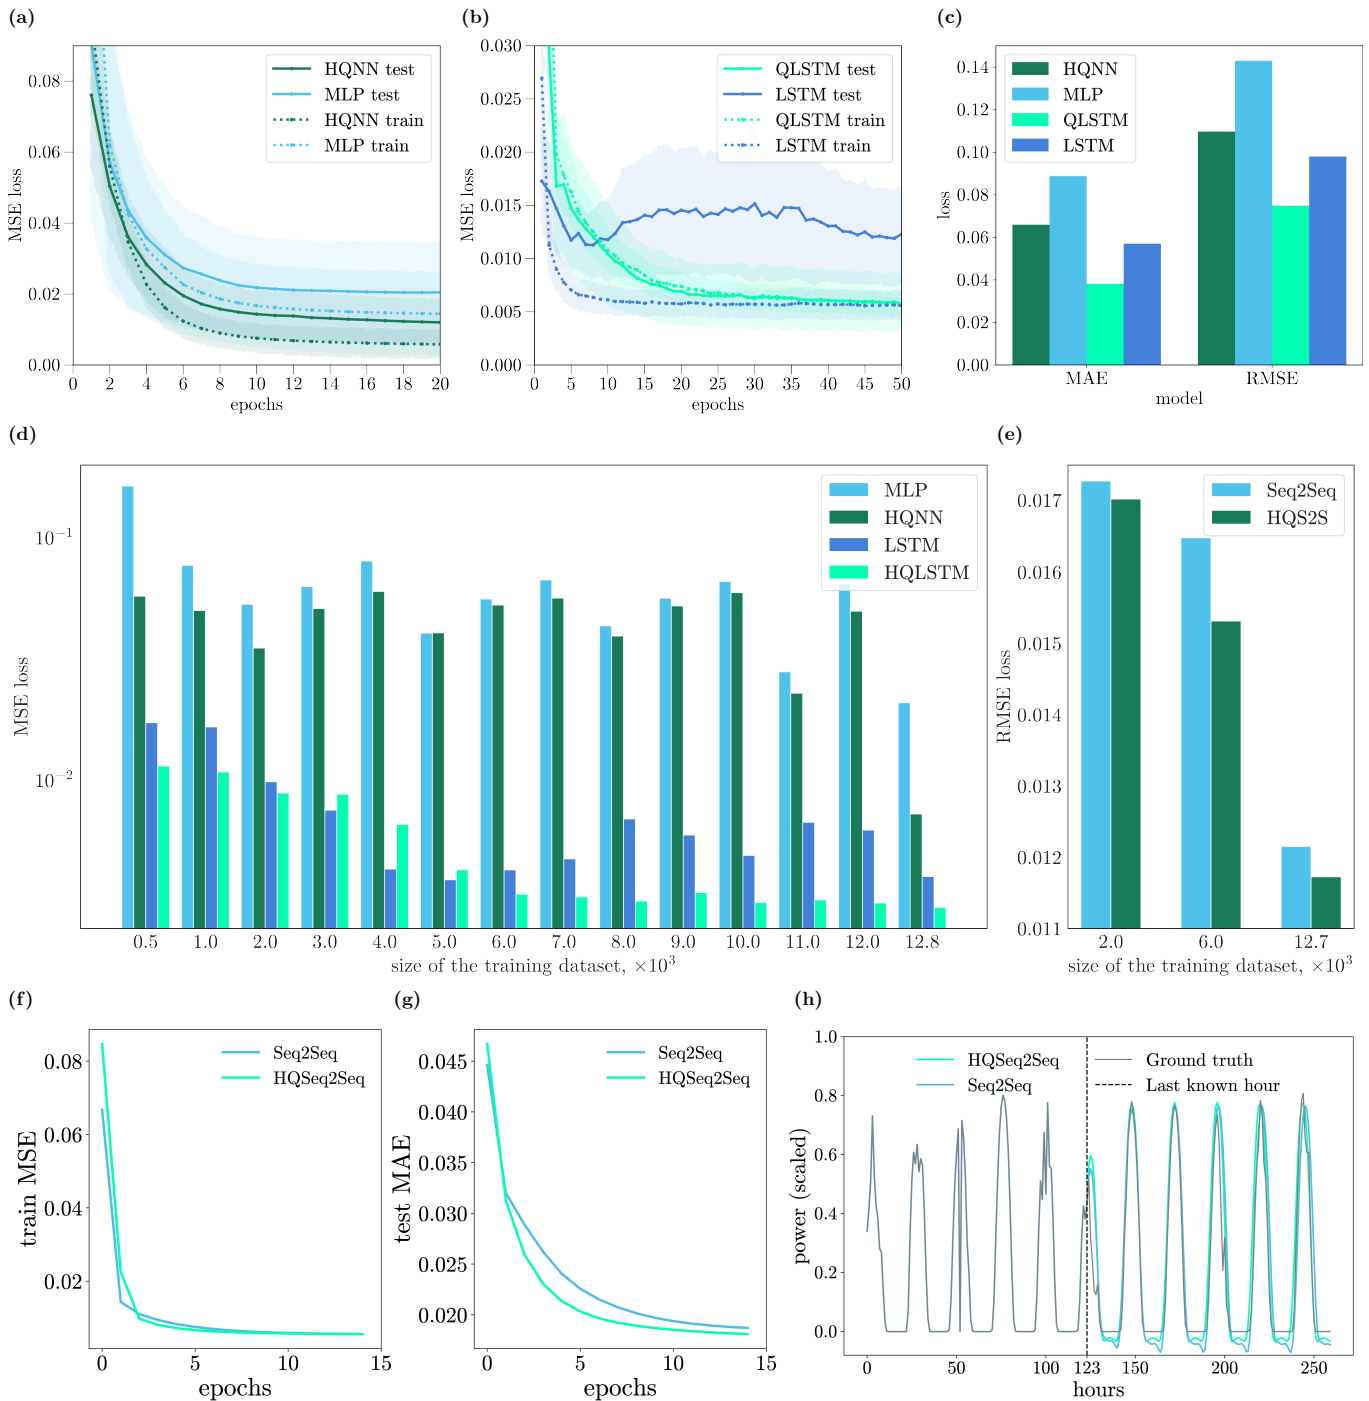


FIG. 4: Results of training and testing of HQNN & MLP and HQLSTM & LSTM models. (a-b) Training and testing history of models. Solid and dotted lines represent training and testing of models respectively. Filled space shows standard deviation of models that were averaged over different testing subsets. Histogram for reduced training dataset. (c) Histogram shows mean MAE and RMSE metrics for our models on testing subset, averaged over different testing subsets. (f-g) Train and test learning curves for the Seq2Seq and HQSeq2Seq models. (h) Example of the classical and hybrid Seq2Seq models inference on the testing data. The models get 124 hours of data as an input (before the dashed line) and gives forecast of “Power” into 137 hours ahead (after the dashed line). The solid black line represents the ground truth value of the “Power” feature. MSE/MAE errors in this particular example are 0.0052/0.0349 for Seq2Seq and 0.0040/0.0292 for HQSeq2Seq.

qubit devices. For our requirements, we selected the lightning.qubit device which is a custom backend for sim-

ulating quantum state-vector evolution. To compute the gradients of the loss function relative to each param-

Hyperparameters	Range	Best value
HQNN		
number of neurons in the second layer	8 – 128	17
number of qubits	2 – 10	8
number of variational layers	1 – 10	7
embedding	R_x, R_y, R_z	R_x
measurement	X, Y, Z	Z
variational part	basic, strongly	basic
initial learning rate	$1 - 1000 \times 10^{-4}$	3×10^{-2}
MLP		
number of neurons in the first layer	8 – 128	32
number of neurons in the second layer	8 – 128	3
number of neurons in the third layer	8 – 128	3
initial learning rate	$1 - 1000 \times 10^{-4}$	1×10^{-2}
HQLSTM		
number of neurons in the second layer	3 – 25	20
dropout range	0 – 99	0.239
number of qubits	2 – 5	4
number of variational layers	1 – 5	1
number of quantum layers	1 – 5	3
initial learning rate	$1 - 100 \times 10^{-3}$	0.52×10^{-2}
LSTM		
number of neurons in the second layer	3 – 25	21
dropout range	0 – 99	0.158
initial learning rate	$1 - 100 \times 10^{-3}$	0.5×10^{-2}

TABLE I: The table shows which hyperparameters are being optimized, the limits of change, and the best values found during hyperparameter optimization.

ter: for the classical components of our hybrid models, we employed the widely-recognized backpropagation algorithm [62]; and for the quantum part, we used the adjoint method as highlighted in Refs. [63, 64].

1. HQNN & MLP and HQLSTM & LSTM

Both the HQNN and its classical analog, the MLP, were trained for 20 epochs. In contrast, the HQLSTM and its classical counterpart, the LSTM, were trained for over 50 epochs. The Adam optimiser [65] from the PyTorch framework was used to update the parameters of the models in order to minimize their loss functions. The comprehensive training process, accompanied by the results, is delineated in Fig. 4 and in Table II.

In this study, we employed cross-validation as a fundamental technique to assess the performance of our models across distinct testing subsets. The application of cross-validation is pivotal to safeguard against potential data leakage from the training dataset into the testing dataset. To achieve this, a rigorous approach was adopted wherein a 24-hour time window, from each side of the subsets, was

Model	HQNN	MLP	HQLSTM	LSTM
train, MAE	0.0458	0.0651	0.0382	0.0454
test, MAE	0.0659	0.0887	0.0343	0.0570
train, MSE	0.0059	0.0144	0.0056	0.0054
test, MSE	0.0120	0.0204	0.0058	0.0096
test, RMSE	0.1097	0.1428	0.0743	0.0937

TABLE II: Summary of the results for the proposed models. In a direct comparison, the HQNN outperforms the MLP in both training and testing losses, as evidenced across three critical metrics: RMSE, MAE, and MSE. Notably, the HQNN’s MSE is 41% less than MLP’s one, which means that HQNN has more reliable forecasts, despite the fact that it has 1.8 times fewer parameters. At the same time, the quality of HQLSTM’s on the test dataset is 40% better compared to LSTM’s one on MAE and MSE metrics, and is 21% better on RMSE metric, although the first one has less than half weights. In a broader comparison encompassing all four models, HQLSTM emerges as the most precise model on all metrics, namely on 52% more precise than HQNN having two times fewer trainable parameters.

systematically excluded from the dataset.

Furthermore, we opted to partition the dataset into training and testing sets in a 4 : 1 ratio. This strategy was implemented to promote a comprehensive evaluation of our models, as we carried out model training and assessment on five distinct data splits. Subsequently, a meticulous averaging process was employed to consolidate the results obtained from these splits, and the mean values thus derived served as the primary metrics for inter-model comparisons.

The utilization of cross-validation techniques in our methodology significantly bolsters the robustness and reliability of our results, as they diminish the reliance on specific train-test partitioning, thereby enhancing the credibility of our findings.

In a head-to-head comparison between HQNN and MLP, the former exhibits superior performance regarding training and testing losses across three key metrics: RMSE, MAE, and MSE. Specifically, HQNN surpasses MLP’s power prediction accuracy by 41% estimated by MSE loss function and by 26% by MAE loss, all the while boasting 1.8 times fewer parameters (2266 & 3987).

On juxtaposing HQLSTM with LSTM, the former outperforms in training and testing loss across all three aforementioned metrics. Remarkably, HQLSTM’s has better predictive ability (on 40% better) than LSTM assessed by the MSE and MAE metrics, and it achieves this with less than half the number of parameters (1109 & 2857). Moreover, HQLSTMs are more resistant to overfitting, while classical LSTM suffers from it.

In a broader comparison encompassing all four models, HQLSTM emerges as the most precise model on all metrics, namely on 52% more precise than HQNN, having two times fewer trainable parameters.

It is worth noting that we performed hyperparameter optimization technique using the Optuna optimizer [66]. The set of optimized parameters, limits of their variation,

and best sets of hyperparameters for all our 4 models are presented in Table I.

Moreover, we scanned external articles that refer to this dataset and found only one article that solves 1 hour ahead PV power prediction using neural networks. In comparison, our HQNN model is better than the model from the external article according to VAF metrics by 40% (91 & 65).

Further, to confirm that hybrid models train better including on a smaller dataset, an additional experiment was conducted, wherein the volume of training data was intentionally reduced. The results are shown in Figure 4. The hybrid models performed better with less data, having fewer losses and better prediction capabilities than classical models.

2. HQSeq2Seq & Seq2Seq

After preprocessing the dataset, which is described in the Section II A, it spanned 12775 hours from 3/5/12 4:55 AM to 8/19/13 10:00 AM for training, and 3194 hours from 8/19/13 11:00 PM to 12/30/13 00:00 AM for testing.

Although models are capable of being trained on sequences of arbitrary lengths, we chose to use sequences of fixed 96 hours for simplicity. In this case, the encoder gets 96 hours of all available features, while the decoder is asked to extrapolate only the “Power” feature of the data 96 hours ahead. Training for 15 epochs with the Adam optimizer (learning rate 0.001) proved to be enough for the models to converge (Fig. 4 (f-g)). As an example of inference, we pass the time series of the length different from 96 into both models and prompt them to give us a forecast for 137 hours ahead (Fig. 4 (h)). We can conclude that both models transition from the fixed sequence length to an arbitrary one quite well. It may even be possible to improve these results by introducing variable-length sequences into the training stage. We also measured the dependency of test loss on the size of the training dataset for Seq2Seq and HQSeq2Seq, shown in Fig. 4 (e). As one can see, the test RMSE loss of the hybrid model is less for any size of training data, which proves that the hybrid model shows an advantage over the classical model, including on a trimmed dataset.

III. DISCUSSION

In this work we introduced three hybrid quantum approaches to time series prediction task. The first two models allow one to predict the power of solar panels for 1 hour ahead, using weather features for the previous 24 hours. The third model allows to predict a longer-term user-defined forecast, showcasing the versatility of our models for various planning tasks.

The first approach is the HQNN, a combination of classical fully connected layers and quantum layer, which is

a VVRQ, analog of classical fully-connected layer. We compared this hybrid model with its classical counterpart, MLP, and demonstrated that, even though HQNN has 1.8 times less variational parameters, it has 41% better predictive ability, estimated by MSE error.

The second approach is HQLSTM, a hybrid quantum analogue of classical LSTM. Here, QDI layer is inserted into each gate of the LSTM cell. This approach provides a 40% improvement in prediction using the MAE and MSE metrics compared to its classical counterpart. Our proposed architecture is a unique combination of classical and quantum layers, which we believe to be a breakthrough in solving time series prediction tasks. Comparing HQNN and HQLSTM models, the second one was better on 52% than HQNN having two times fewer weights.

The third approach is hybrid Seq2Seq model, a classical Seq2Seq model, consisting of 2 LSTMs with quantum layer at the end. This approach allows one to predict the PV power not only for an hour ahead, but for any number of hours ahead, without knowing the weather features in advance. The addition of the proposed QDI layer improves the accuracy of the predictions, reducing the MAE error by 16% compared to a purely classical Seq2Seq model. Our proposed architecture is a unique combination of classical and quantum layers, which we believe to be a breakthrough in solving time series prediction tasks.

Also, for all models, we conducted an additional experiment in which our models were trained on a reduced dataset, and confirmed that hybrid models have better learning capabilities and have less loss trained on any amount of dataset compared to their classical counterparts. This confirmation can serve as an excellent motivation to use hybrid networks for applications where data collection is a complex task.

It is worth noting that all parameters are trainable in our layers; the architecture and hyperparameters were selected by the Optuna optimizer. Moreover, we compare our models to a paper that solved the same problem using the same dataset, and demonstrate that our best HQLSTM is 40% more accurate in predicting power using VAF metric.

To fully unlock the potential of HQNNs in time series prediction problem, further research and testing of models on other datasets is necessary. Also, the development of more efficient optimization VQC training and implementation methods, larger-scale quantum hardware could lead to even more significant performance improvements.

Furthermore, while this work was done on a public dataset with an emphasis on hybrid quantum models for better forecasting performance, the quality and source of data plays a crucial role in overall effectiveness in the real world, especially considering weather data. Accurate weather forecast is a crucial input into any high performing and useful PV prediction given its dynamism and influence on PV output. An interesting area of re-

search is cloud prediction using satellite and weather data for geo-locations, directly impacting solar irradiance and therefore PV output. The added complexity could enhance the need for hybrid quantum models to increase computational efficiency and higher quality forecasts.

To summarize, our developments provide three hybrid quantum approaches for time series problems that

demonstrate the possibility of combining classical and quantum methods. Our proposed models show improved performance compared to classical models with similar architecture when using fewer variation parameters. We believe that these results pave the way for further research in developing hybrid models that leverage the strengths of both classical and quantum computing.

-
- [1] M. Sabri and M. El Hassouni. A novel deep learning approach for short term photovoltaic power forecasting based on gru-cnn model. *E3S Web of Conferences*, 336:00064, 2022.
- [2] B. Li and J. Zhang. A review on the integration of probabilistic solar forecasting in power systems. *Solar Energy*, 210:68–86, 2020.
- [3] O. Ellabban, H. Abu-Rub, and F. Blaabjerg. Renewable energy resources: Current status, future prospects and their enabling technology. *Renewable and Sustainable Energy Reviews*, 39:748–764, 2014.
- [4] Anthony R. Florita a Carlo Brancucci Martinez-Anido a, Benjamin Botor a. The value of day-ahead solar power forecasting improvement. *Science*, 129:192–203, 2016.
- [5] D. Lee and K. Kim. PV power prediction in a peak zone using recurrent neural networks in the absence of future meteorological information. *Renewable Energy*, 173:1098–1110, 2021.
- [6] S.A. Bozorgavari, J. Aghaei, S. Pirouzi, V. Vahidinassab, H. Farahmand, and M. Korpás. Two-stage hybrid stochastic/robust optimal coordination of distributed battery storage planning and flexible energy management in smart distribution network. *Journal of Energy Storage*, 26:100970, 2019.
- [7] M. Diagne, M. David, P. Lauret, J. Boland, and N. Schmutz. Review of solar irradiance forecasting methods and a proposition for small-scale insular grids. *Renewable and Sustainable Energy Reviews*, 27:65–76, 2013.
- [8] A. Ahmed and M. Khalid. A review on the selected applications of forecasting models in renewable power systems. *Renewable and Sustainable Energy Reviews*, 100:9–21, 2019.
- [9] J. López Gómez, A. Ogando Martínez, F. Troncoso Pastoriza, L. Febrero Garrido, E. Granada Álvarez, and J. A. Orosa García. Photovoltaic Power Prediction Using Artificial Neural Networks and Numerical Weather Data. *Sustainability*, 12:10295, 2020.
- [10] M. R. Kaloop, A. Bardhan, N. Kardani, P. Samui, J. W. Hu, and A. Ramzy. Novel application of adaptive swarm intelligence techniques coupled with adaptive network-based fuzzy inference system in predicting photovoltaic power. *Renewable and Sustainable Energy Reviews*, 148:111315, 2021.
- [11] R. S. Kulkarni, D. B. Talange, and N. V. Mate. Output Estimation of Solar Photovoltaic (PV) System. In *2018 International Symposium on Advanced Electrical and Communication Technologies (ISAECT)*, pages 1–6. IEEE, 2018.
- [12] J. Shi, W.J. Lee, Y. Liu, Y. Yang, and P. Wang. Forecasting Power Output of Photovoltaic Systems Based on Weather Classification and Support Vector Machines. *IEEE Transactions on Industry Applications*, 48(3):1064–1069, 2012.
- [13] O. P. Rocha, A.M. da Silva1, and A. Á. B. Santos. Computational Model for Photovoltaic Solar Energy Forecasting Based on the K-Nearest Neighbor Method. *Journal of Bioengineering, Technologies and Health*, 5(3):168–172, 2022.
- [14] R. H. Inman, H. T. C. Pedro, and C. F. M. Coimbra. Solar forecasting methods for renewable energy integration. *Progress in Energy and Combustion Science*, 39(6):535–576, 2013.
- [15] R. Blaga, A. Sabadus, N. Stefu, C. Dughir, M. Paulescu, and V. Badescu. A current perspective on the accuracy of incoming solar energy forecasting. *Progress in energy and combustion science*, 70:119–144, 2019.
- [16] M. Schuld, I. Sinayskiy, and F. Petruccione. An introduction to quantum machine learning. *Contemporary Physics*, 56(2):172–185, 2015.
- [17] J. Biamonte, P. Wittek, N. Pancotti, P. Rebentrost, N. Wiebe, and S. Lloyd. Quantum machine learning. *Nature*, 549(7671):195–202, 2017.
- [18] V. Dunjko and H. J Briegel. Machine learning & artificial intelligence in the quantum domain: a review of recent progress. *Reports on Progress in Physics*, 81(7):074001, 2018.
- [19] Alexey Melnikov, Mohammad Kordzanganeh, Alexander Alodjants, and Ray-Kuang Lee. Quantum machine learning: from physics to software engineering. *Advances in Physics: X*, 8(1):2165452, 2023.
- [20] D. Emmanoulopoulos and S. Dimoska. Quantum Machine Learning in Finance: Time Series Forecasting. arXiv preprint arXiv:2202.00599, 2019.
- [21] M. M. Sushmit and I. M. Mahbulbul. Forecasting solar irradiance with hybrid classical-quantum models: A comprehensive evaluation of deep learning and quantum-enhanced techniques. *Energy Conversion and Management*, 294:117555, 2023.
- [22] M. Schuld, M. Fingerhuth, and F. Petruccione. Implementing a distance-based classifier with a quantum interference circuit. *Europhysics Letters*, 119(6):60002, 2018.
- [23] S. Lloyd, M. Mohseni, and P. Rebentrost. Quantum algorithms for supervised and unsupervised machine learning. arXiv preprint arXiv:1307.0411, 2013.
- [24] P.W. Shor. Algorithms for quantum computation: discrete logarithms and factoring. *Proceedings 35th Annual Symposium on Foundations of Computer Science*, pages 124–134, 1994.
- [25] S. Lloyd. Universal quantum simulators. *Science*, 273:1073–1078, 1996.
- [26] M. Schuld, A. Bocharov, K. M Svore, and N. Wiebe. Circuit-centric quantum classifiers. *Physical Review A*, 101(3):032308, 2020.

- [27] S. Aaronson and L. Chen. Complexity-Theoretic Foundations of Quantum Supremacy Experiments. [arXiv preprint arXiv:1612.05903](#), 2016.
- [28] P. Jain and S. Ganguly. Hybrid quantum generative adversarial networks for molecular simulation and drug discovery. [arXiv preprint arXiv:2212.07826](#), 2022.
- [29] Asel Sagingalieva, Mohammad Kordzanganeh, Nurbolat Kenbayev, Daria Kosichkina, Tatiana Tomashuk, and Alexey Melnikov. Hybrid quantum neural network for drug response prediction. *Cancers*, 15(10):2705, 2023.
- [30] Luca Lusnig, Asel Sagingalieva, Mikhail Surmach, Tatjana Protasevich, Ovidiu Michiu, Joseph McLoughlin, Christopher Mansell, Graziano de’ Petris, Deborah Bonazza, Fabrizio Zanconati, Alexey Melnikov, and Fabio Cavalli. Hybrid quantum image classification and federated learning for hepatic steatosis diagnosis. [arXiv preprint arXiv:2311.02402](#), 2023.
- [31] A. Kurkin, J. Hegemann, M. Kordzanganeh, and A. Melnikov. Forecasting the steam mass flow in a power-plant using the parallel hybrid network. [arXiv preprint arXiv:2307.09483](#), 2023.
- [32] Serge Rainjonneau, Igor Tokarev, Sergei Iudin, Saaketh Rayaprolu, Karan Pinto, Daria Lemtiuzhnikova, Miras Koblan, Egor Barashov, Mo Kordzanganeh, Markus Pflitsch, and Alexey Melnikov. Quantum algorithms applied to satellite mission planning for Earth observation. *IEEE Journal of Selected Topics in Applied Earth Observations and Remote Sensing*, 16:7062–7075, 2023.
- [33] Nathan Haboury, Mo Kordzanganeh, Sebastian Schmitt, Ayush Joshi, Igor Tokarev, Lukas Abdallah, Andrii Kurkin, Basil Kyriacou, and Alexey Melnikov. A supervised hybrid quantum machine learning solution to the emergency escape routing problem. [arXiv preprint arXiv:2307.15682](#), 2023.
- [34] Asel Sagingalieva, Andrii Kurkin, Artem Melnikov, Daniil Kuhmistrov, et al. Hybrid quantum ResNet for car classification and its hyperparameter optimization. *Quantum Machine Intelligence*, 5(2):38, 2023.
- [35] Mo Kordzanganeh, Pavel Sekatski, Leonid Fedichkin, and Alexey Melnikov. An exponentially-growing family of universal quantum circuits. *Machine Learning: Science and Technology*, 4(3):035036, 2023.
- [36] Arsenii Senokosov, Alexander Sedykh, Asel Sagingalieva, and Alexey Melnikov. Quantum machine learning for image classification. [arXiv preprint arXiv:2304.09224](#), 2023.
- [37] Michael Perelshtein, Asel Sagingalieva, Karan Pinto, Vishal Shete, Alexey Pakhomchik, Artem Melnikov, et al. Practical application-specific advantage through hybrid quantum computing. [arXiv preprint arXiv:2205.04858](#), 2022.
- [38] M. Schuld, R. Sweke, and J. J. Meyer. Effect of data encoding on the expressive power of variational quantum-machine-learning models. *Physical Review A*, 103(3):032430, 2021.
- [39] Mo Kordzanganeh, Daria Kosichkina, and Alexey Melnikov. Parallel hybrid networks: an interplay between quantum and classical neural networks. *Intelligent Computing*, 2:0028, 2023.
- [40] Alexandr Sedykh, Maninadh Podapaka, Asel Sagingalieva, Nikita Smertyak, Karan Pinto, Markus Pflitsch, and Alexey Melnikov. Quantum physics-informed neural networks for simulating computational fluid dynamics in complex shapes. [arXiv preprint arXiv:2304.11247](#), 2023.
- [41] M Malvoni, M.G. De Giorgi, and P.M. Congedo. Data on photovoltaic power forecasting models for Mediterranean climate. *Data in Brief*, 7:1639–1642, 2016.
- [42] J. Aghaei, S.A. Bozorgavari, S. Pirouzi, H. Farahmand, and M. Korpås. Flexibility planning of distributed battery energy storage systems in smart distribution networks. *Iranian Journal of Science and Technology, Transactions of Electrical Engineering*, 44(3):1105–1121, 2019.
- [43] Rohan S. Kulkarni, Rajani B. Shinde, and Dhananjay B. Talange. Performance Evaluation and a New Thermal Model for a Photovoltaic-Thermal Water Collector System. In *2018 International Conference on Smart Grid and Clean Energy Technologies (ICSGCE)*, pages 106–111. IEEE, 2018.
- [44] M. Malvoni, M. G. De Giorgi, and P. M. Congedo. Study of degradation of a grid connected photovoltaic system. *Energy Procedia*, 126:644–650, 2017.
- [45] F. Bloch. Nuclear Induction. *Physical Review*, 70(7-8):460–474, 1946.
- [46] Adriano Barenco, Charles H. Bennett, Richard Cleve, David P. DiVincenzo, Norman Margolus, Peter Shor, Tycho Sleator, John A. Smolin, and Harald Weinfurter. Elementary gates for quantum computation. *Physical Review A*, 52(5):3457–3467, 1995.
- [47] S. Hochreiter and J. Schmidhuber. Long short-term memory. *Neural computation*, 9(8):1735–1780, 1997.
- [48] H. Gao, S. Qiu, J. Fang, N. Ma, J. Wang, K. Cheng, et al. Short-Term Prediction of PV Power Based on Combined Modal Decomposition and NARX-LSTM-LightGBM. *Sustainability*, 15(10):8266, 2023.
- [49] H. Sharadga, S. Hajimirza, and R. S. Balog. Time series forecasting of solar power generation for large-scale photovoltaic plants. *Renewable Energy*, 150:797–807, 2020.
- [50] M. Tovar, M. Robles, and F. Rashid. PV Power Prediction, Using CNN-LSTM Hybrid Neural Network Model. Case of Study: Temixco-Morelos, México. *Energies*, 13(24):6512, 2020.
- [51] X. Wang, X. Wang, and S. Zhang. Adverse Drug Reaction Detection from Social Media Based on Quantum Bi-LSTM with Attention. *IEEE Access*, 11:16194–16202, 2022.
- [52] M. S. Akter, H. Shahriar, and Z. A. Bhuiya. Automated vulnerability detection in source code using quantum natural language processing. In *International Conference on Ubiquitous Security*, pages 83–102. Springer, 2022.
- [53] Y. Yu, G. Hu, C. Liu, J. Xiong, and Z. Wu. Prediction of solar irradiance one hour ahead based on quantum long short-term memory network. *IEEE Transactions on Quantum Engineering*, 2023.
- [54] S. Y.C. Chen, S. Yoo, and Y. L. L. Fang. Quantum long short-term memory. [arXiv preprint arXiv:2009.01783](#), 2020.
- [55] I. Sutskever, O. Vinyals, and Q. V Le. Sequence to sequence learning with neural networks. *Advances in neural information processing systems*, 27, 2014.
- [56] N. Kalchbrenner and P. Blunsom. Recurrent continuous translation models. In *Proceedings of the 2013 conference on empirical methods in natural language processing*, pages 1700–1709, 2013.
- [57] Y. Mu, M. Wang, X. Zheng, and H. Gao. An improved LSTM-Seq2Seq-based forecasting method for electricity load. *Frontiers in Energy Research*, 10:1093667, 2023.
- [58] Adrián Pérez-Salinas, Alba Cervera-Lierta, Elies Gil-Fuster, and José I. Latorre. Data re-uploading for a uni-

- versal quantum classifier. *Quantum*, 4:226, 2020.
- [59] QMware. QMware — The first global quantum cloud. <https://qm-ware.com/>, 2022.
- [60] Mohammad Kordzanganeh, Markus Buchberger, Basil Kyriacou, Maxim Povolotskii, Wilhelm Fischer, Andrii Kurkin, et al. Benchmarking simulated and physical quantum processing units using quantum and hybrid algorithms. *Advanced Quantum Technologies*, 6(8):2300043, 2023.
- [61] PyTorch — Machine Learning framework. <https://pytorch.org/>, 2022.
- [62] D. E. Rumelhart, E. Hinton, G, and R. J. Williams. Learning representations by back-propagating errors. *Nature*, 323(6088):533–536, 10 1986.
- [63] X. Z. Luo, J. G. Liu, P. Zhang, and L. Wang. Yao.jl: Extensible, Efficient Framework for Quantum Algorithm Design. *Quantum*, 4:341, 2020.
- [64] T. Jones and J. Gacon. Efficient calculation of gradients in classical simulations of variational quantum algorithms. *arXiv preprint arXiv:2009.02823*, 2020.
- [65] Diederik P. Kingma and Jimmy Ba. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*, 2017.
- [66] Optuna — A hyperparameter optimization framework. <https://optuna.org/>, 2023.