

Review

Not peer-reviewed version

Quantum Machine Learning Early Opportunities for the Energy Industry: a Systematic Review

[Francesco Strata](#)^{*}, Luca Migliori, Nour Gebran, Nicolina Guarino, Giacomo Carlo Colombo, Sara Pezzuolo, Emiliano Luzietti

Posted Date: 24 April 2025

doi: 10.20944/preprints202504.2070.v1

Keywords: quantum computing; innovation; energy sector; global energy demand; sustainability; computational technologies; energy production; energy storage; energy distribution; Artificial Intelligence (AI); Machine Learning (ML); Quantum Machine Learning (QML); systematic review; energy industry value chain; use cases; feasibility; benefits; Assessment Model for Innovation Management (AMIM); technological maturity; market compatibility



Preprints.org is a free multidisciplinary platform providing preprint service that is dedicated to making early versions of research outputs permanently available and citable. Preprints posted at Preprints.org appear in Web of Science, Crossref, Google Scholar, Scilit, Europe PMC.

Copyright: This open access article is published under a Creative Commons CC BY 4.0 license, which permit the free download, distribution, and reuse, provided that the author and preprint are cited in any reuse.

Article

Quantum Machine Learning Early Opportunities for the Energy Industry: a Systematic Review

Francesco Strata ^{1,*}, Luca Migliori ¹, Nour Gebran ¹, Nicolina Guarino ², Giacomo Carlo Colombo ¹, Sara Pezzuolo ¹ and Emiliano Luzietti ¹

¹ Data & AI, PricewaterhouseCoopers, Milan, Italy

² Data & AI, PricewaterhouseCoopers, Napoli, Italy

* Correspondence: francesco.strata@pwc.com

Abstract: Quantum computing innovations have garnered significant attention for their potential to revolutionize industries, with the energy sector being one of the most promising areas for application. As global energy demand increases and sustainability becomes more critical, computational technologies offer groundbreaking solutions for energy production, storage, and distribution. In this landscape, quantum computing plays a crucial role in unlocking the full potential of artificial intelligence (AI) and machine learning (ML), as the research and development in the quantum machine learning (QML) field grows constantly. In this paper, we present a systematic review of early quantum machine learning applications within the energy industry value chain. We analyze and discuss various use cases in the energy sector, thoroughly examining each one to understand its potential applications and impact. We then evaluate these early-stage quantum applications to determine their feasibility and benefits, offering insights into their relevance and effectiveness in the context of the industry's evolving landscape. This is done by introducing a novel framework: the Assessment Model for Innovation Management (AMIM). Our research highlights the opportunities quantum innovations present for the energy sector and offers actionable insights into which applications are the best investments and why. Overall, the feasibility and technological maturity of QML use cases are still in the early stages, though their market compatibility and potential benefits are relatively high for most of them. This indicates that while quantum machine learning holds immense potential, further development is necessary to fully realize its benefits in the energy sector.

Keywords: quantum computing; innovation; energy sector; global energy demand; sustainability; computational technologies; energy production; energy storage; energy distribution; Artificial Intelligence (AI); Machine Learning (ML); Quantum Machine Learning (QML); systematic review; energy industry value chain; use cases; feasibility; benefits; Assessment Model for Innovation Management (AMIM); technological maturity; market compatibility

1. Introduction

Quantum Computing is a very promising and potentially impactful technology, given its theoretical capacity to outperform classical computing in a variety of complex computational workloads [34,47,126]. In this view, a multidisciplinary ecosystem thrives on academia and industry, looking for early applications of quantum computing able to solve practical challenges in optimization, simulation, machine learning (ML) and others. In particular, QC's features nicely suit challenging mathematical problems that are crucial to machine learning algorithms, such as classification, forecasting, clustering, and others [10]. Given that ML is more and more pervasive in modern industry, the field of quantum machine learning (QML) raises a great interest for the foreseen added value that it could bring to strategic data-driven applications.

Nowadays more than ever, Energy and Utilities companies are facing hard challenges: the increasing complexity of power grids, delocalized and volatile energy sources, the ever increasing energy-demand, decarbonization objectives, free market trading, etc. Among others, machine learning

routines are critical to many power supply operations such as demand forecasting, RE generation forecasting, fault diagnosis and grid stability assessment. Plus, energy market management is crucial for the financial stability of power supply companies, as well as sustainable end-users offerings.

Although there are no commercially available quantum workloads yet, the rapid growth of the technology and the new design of hybrid algorithms has already allowed some experimental works that could be pivotal for short-term applications. Interestingly, these works try to find an empirical and actually prospective quantum advantage over the hardware or emulators currently available. In the following article they will be reviewed, particularly the ones that are concerned with the energy sector. In the first chapter the fundamental concepts of quantum computing will be presented. In the second chapter, the main machine learning techniques used in the considered industrial applications are reported. In the third chapter, an overview of the field of quantum machine learning is given, with a particular focus on the variational approach, which constitutes the hope of using quantum technologies in the short to medium term. In the fourth chapter, the applications of early-stage QML in the different use cases of the energy sector are reviewed. Finally, the Assessment Model for Innovative Management (AMIM) is developed and the various use cases are evaluated.

2. Quantum Computing Fundamentals

A quantum computer is a universal computing device that uses quantum bits, or qubits, to store information and execute computations by harnessing the distinctive characteristics of quantum mechanics. This type of computing, known as quantum computing, involves gathering various states of qubits, including superposition, interference, and entanglement, to carry out calculations [48].

In classical computers, the elementary unit of information storage is a bit, which can exist in one of two states: 0 or 1. In contrast, the elementary unit of quantum information is the qubit, which can exist in states represented as $|0\rangle$ and $|1\rangle$, or in a superposition of these states, denoted as $|\psi\rangle$. The states $|0\rangle$ and $|1\rangle$ are referred to as computational basis states, and the Dirac notation " $|\cdot\rangle$ " is the standard representation for qubit states in quantum mechanics. A quantum system with n qubits has a Hilbert space of 2^n dimensions, allowing for 2^n mutually orthogonal quantum states. For instance, with three qubits, the states can be represented as $|000\rangle, |001\rangle, |010\rangle, |011\rangle, |100\rangle, |101\rangle, |110\rangle, |111\rangle$.

One of the key features of qubits is *superposition*, which means that a qubit can be in a combination of both $|0\rangle$ and $|1\rangle$ states simultaneously, described by probability amplitudes α and β . Specifically, we can represent a qubit as follow:

$$|\psi\rangle = \alpha|0\rangle + \beta|1\rangle \quad (1)$$

where α and β are probability amplitudes of the basis states $|0\rangle$ and $|1\rangle$, respectively, and both α and β are complex numbers and satisfy $|\alpha|^2 + |\beta|^2 = 1$. Upon measurement, the qubit's state collapses to either $|0\rangle$ or $|1\rangle$ based on the probability of $|\alpha|^2$ or $|\beta|^2$, respectively.

Another feature in qubits is *entanglement*, which describes a strong correlation between qubits, regardless of the distance separating them. When two qubits are entangled, a change to one qubit will instantaneously affect the other, even if they are far apart. Moreover, when one entangled qubit is measured, its result contains the information of the other entangled qubits.

By exploiting both features, superposition and entanglement, quantum computers can outperform classical computers in terms of computing power and speed.

The approaches to physically represent and manipulate qubits can be categorized into two main types: analog quantum computing and digital gate-based quantum computing as discussed in the following paragraph.

2.1. Types of Hardware

To apply quantum computing, types of computing systems should be chosen between quantum gate-based model and quantum annealing model:

Analog quantum computing In analog quantum computing, the quantum state is smoothly changed by quantum operations such that the information encoded in the final system corresponds to the desired answer with high probability. One example of analog quantum computing is the adiabatic quantum computer [5], which refer to the idea of building a type of universal quantum computing. A special form of adiabatic quantum computers is quantum annealing, which is a framework that incorporates algorithms and hardware designed to solve computational problems via adiabatic quantum evolution towards the ground states of a given quantum system [136]. Quantum annealing takes advantage of the fact that physical systems strive towards the state with the lowest energy, e.g. hot things cool down over time or objects roll downhill. As such, in quantum annealing the energetically most favorable state then corresponds to the solution of the optimization problem [5]. Using the property of superposition and entanglement, the quantum annealer is able to calculate all potential solutions at the same time, which speeds up the calculation process drastically in comparison to classical computers [124]. Compared to the quantum circuit model, the quantum annealing model is so far less susceptible to noise. Thus, quantum annealing can utilize more qubits and have greater scalability than the quantum gate-based approach. Conversely, quantum annealing is a special-purpose paradigm for optimization problems or probabilistic sampling and no theoretically proven computational advantage has been shown yet.

Digital gate-based quantum computing In digital gate-based quantum computing, the information encoded in qubits is manipulated through digital gates. In comparison to the analog approach in which you sample the natural evolution of quantum states to find the optimal state of low energy, in digital gate-based quantum computers the evolution of the quantum states is manipulated in terms of activity and controlled to find the optimal solution [29]. By actively manipulating the state of qubits, gate-based computing gains a significant advantage in flexibility, enabling such a processing unit to solve a broader range of problems compared to quantum annealing. Indeed, digital gate-based quantum computing is conceptually very similar to classical computation [48]. While a classical algorithm is running on a computer as a series of instructions (gates such as AND, OR, NOT, ...), quantum gate-based calculators manipulate individual or pairs of classical bits and flip them between zero and one states according to a set of rules. Quantum gates operate directly on one or multiple qubits by rotating and shifting them between different superpositions of the zero and one states as well as different entangled states.

2.2. Quantum Error: Error Suppression, Error Mitigation and Error Correction

In quantum computing, there are many potential sources of noise. For example, quantum computers, especially digital gate-based ones, are highly sensitive to environmental changes such as vibrations and temperature fluctuations. Noise can also arise from imprecise control of the quantum hardware or manufacturing defects[29]. Most quantum computers require their chips to be cooled to a hundredth of a degree above absolute zero to operate. Since noise cannot be avoided, the first era of quantum computers is called the *Noisy Intermediate-Scale Quantum (NISQ) era* [104]. This term implies that current quantum hardware, which uses dozens of qubits, has high error rates that need improvement before we can build useful quantum computers with hundreds or even thousands of usable qubits. NISQ refers to a class of quantum devices that operate with tens to a few hundred qubits and are subject to significant levels of quantum noise and error. These devices represent the current frontier of quantum computing technology, bridging the gap between small-scale experimental systems and future large-scale, fault-tolerant quantum computers. NISQ devices are characterized by their intermediate size and the presence of noise in their operations. While they are large enough to perform non-trivial quantum computations, they lack the full error correction capabilities required for fault-tolerant computing. This makes them susceptible to errors and limits the complexity and accuracy of the computations they can perform. However, NISQ devices are still powerful tools for exploring quantum algorithms and applications.

Quantum computers today have high error rates – around 1 error occurs in 1000 operations before failure. For quantum computers to be useful, error rates need to be as low as 1 in a trillion. A huge

improvement in performance is needed, and proven progress is already happening in the community. We can break error handling into three core pieces, each with their own research and development considerations: error suppression, error mitigation, and error correction:

- *Error suppression* is a fundamental level of error handling in quantum computing. It involves techniques that use knowledge of undesirable effects to anticipate and avoid potential impacts, often at the hardware level. These methods, which date back decades, typically involve altering or adding control signals to ensure the processor returns the desired result. Error suppression, also known as deterministic or dynamic error suppression, reduces the likelihood of hardware errors during quantum bit manipulation or memory storage. It leverages quantum control techniques to build resilience against errors. For example, quantum logic gates, which are essential for quantum algorithms, can be redefined using machine learning to enhance robustness against errors. Similarly, control operations can protect idle qubits from external interference, akin to a "force field" that deflects noise. Various strategies for error suppression can significantly improve quantum computing performance. Designing new quantum logic gates can make operations up to ten times less likely to suffer errors, thus enhancing algorithmic performance. Research has shown that error suppression can increase the likelihood of achieving correct results by over 1000 times and improve Quantum Volume, a measure of hardware capability. Error suppression can be integrated into quantum firmware or configured for automated workflows, reducing errors on each run without additional overhead. However, it cannot correct all errors, such as "Energy Relaxation" (T1) errors, which require Quantum Error Correction strategies.
- *Error mitigation* (EM) is crucial for making near-term quantum computers useful by reducing or eliminating noise through the estimation of expectation values. Each EM method has its own overhead and accuracy level. The most powerful techniques can have exponential overhead, meaning the time to run increases exponentially with the problem size (number of qubits and circuit depth). Users can choose the best technique based on their accuracy needs and acceptable overhead. In quantum computing, estimating calculated parameters, like energy levels of molecules in quantum chemistry, can be affected by errors in both algorithm execution and measurement. Various strategies have been developed to improve results through post-processing, including randomized compiling, measurement-error mitigation, zero-noise extrapolation, and probabilistic error cancellation. These strategies involve running many slightly different versions of a target algorithm and combining the results to "extract the right answer through the errors". Measurement-error mitigation is particularly powerful, using statistical techniques to identify correct calculations despite readout failures. To maximize benefits from EM, an algorithm might need to be run around 100 times with different configurations, which could lead to a significant increase in quantum computing costs.
- *Error correction* (QEC) aims to achieve fault-tolerant quantum computation by building redundancies so that even if some qubits experience errors, the system still returns accurate results. In classical computing, error correction involves encoding information with redundancy to check for errors. Quantum error correction follows the same principle but must account for new types of errors and carefully measure the system to avoid collapsing the quantum state. In QEC, single qubit values (logical qubits) are encoded across multiple physical qubits. Gates are implemented to treat these physical qubits as error-free logical qubits. The QEC algorithm distributes quantum information across supporting qubits, protecting it against local hardware failures. Special measurements on helper qubits indicate failures without disturbing the stored information, allowing corrections to be applied. QEC involves cycles of gates, syndrome measurements, error inference, and corrections, functioning as feedback stabilization. The entire error-correction cycle is designed to tolerate errors at every stage, enabling error-robust quantum processing even with unreliable components. This fault-tolerant architecture enables the construction of large quantum computers with low error rates, but quantum error correction (QEC) requires a significant number of qubits. The greater the noise, the more qubits are needed, and estimates suggest that thousands

of physical qubits may be required to encode a single protected logical qubit, which presents a challenge given the limited qubit counts of current systems. The sheer scale of this overhead and the complexity of QEC is why despite many promising results, QEC still needs further refinement to provide efficient operations for useful applications. This may change soon though, following the recent advancements from hardware providers.

3. Classical Machine Learning: Principles and Overview

Machine Learning (ML) is a branch of Artificial Intelligence (AI) that focuses on building systems that can learn from and make decisions based on data. Unlike traditional programming, where a computer follows explicit instructions, machine learning involves creating algorithms that allow computers to identify patterns and make predictions or decisions without being explicitly programmed to perform the task. The majority of ML algorithms involve training, which means using large amounts of data and an efficient algorithm to adapt and improve based on the training process. After the training phase, the algorithm will be able to apply the acquired knowledge to situations similar to those encountered during training. The advent of Big Data, along with advancements in Internet of Things (IoT) technologies, has exponentially increased the availability of data for training machine learning (ML) models. Nowadays, ML is used to extract meaningful information from these data in a timely and intelligent way. The insights extracted can be used to build various intelligent applications in relevant domains, for instance, to build a data-driven automated and intelligent cybersecurity system, personalized context-aware smart mobile applications, and so on. Machine Learning algorithms are mainly divided into four categories: Supervised learning, Unsupervised learning, Semi-supervised learning, and Reinforcement learning, according to the amount and type of supervision they get during training.

In supervised learning, the training data you feed to the algorithm includes the desired solutions, called labels. The goal of the algorithm is to learn a mapping from inputs to outputs that can be used to predict the labels of new, unseen data. Supervised learning is commonly used for tasks such as classification and regression. The process involves feeding the algorithm a set of training data, allowing it to adjust its parameters to minimize the error between its predictions and the actual labels. Once trained, the model can make accurate predictions on new data based on the patterns it has learned. A classification algorithm divides inputs into a certain number of categories or classes, based on the labeled data it was trained on. These algorithms can handle both binary classifications (e.g. classifying customer feedback as positive or negative, or emails as spam or not) and multi-class classifications (e.g. recognizing handwritten letters, numbers, or categorizing drugs). This versatility makes classification algorithms useful for a wide range of applications. Regression tasks are distinct because they predict a numerical relationship between the input and output data. These models are designed to forecast continuous values rather than discrete categories. For example, they can estimate real estate prices based on factors like zip codes, property size, and market trends. In digital marketing, regression models can predict click-through rates for online ads by analyzing variables such as the time of day, user demographics, and ad content. Additionally, in consumer behavior analysis, these models can forecast a customer's willingness to pay for a product based on attributes like age, income, and purchasing history.

In unsupervised learning the training data is unlabeled, so the system tries to learn without a teacher, making it a data-driven approach. It is commonly employed to extract generative features, uncover significant patterns and structures, group results, and for exploratory analysis. Typical tasks in unsupervised learning include clustering, density estimation, feature learning, dimensionality reduction, discovering association rules, and detecting anomalies.

3.1. Kernel Method

Kernel methods are a class of techniques used in machine learning to enable algorithms to learn complex patterns in data by implicitly mapping input data into a higher-dimensional space, without explicitly computing the coordinates in that space, a concept known as the "*kernel trick*". This allows

the algorithm to discover relationships that would be difficult or impossible to model in the original feature space and, also, allows for the linear separation of data that is not linearly separable in the original space. From a mathematical point of view, instead of computing the dot product in the high-dimensional space, we use a kernel function to compute the dot product directly in the input space. Common kernel functions are, for example:

- Linear Kernel $K(x, x') = x^T x'$ — This is just the standard dot product in the original space and doesn't map the data to a higher dimension.
- Polynomial Kernel $K(x, x') = (x^T x' + c)^d$ — This maps the data into a higher-dimensional space based on polynomial functions.
- Radial Basis Function (RBF) / Gaussian Kernel $K(x, x') = \exp\left(-\frac{\|x - x'\|^2}{2\sigma^2}\right)$ — This kernel maps the data into an infinite-dimensional space and is often used in SVMs for classification tasks. It is useful for capturing non-linear relationships.
- Sigmoid Kernel $K(x, x') = \tanh(ax^T x' + c)$ — Based on the hyperbolic tangent function, it's similar to the activation function used in neural networks.

Kernel methods can be used for both supervised and unsupervised learning. In supervised learning, kernels allow algorithms like Support Vector Machines to perform classification and regression tasks in a high-dimensional space efficiently. In unsupervised learning, they find applications in methods like kernel spectral clustering, aiding in the identification of inherent groupings within data.

3.2. Random Forest

Random forest is a commonly-used machine learning algorithm that combines the output of multiple decision trees to reach a single result. Its ease of use and flexibility have fueled its adoption, as it handles both classification and regression problems. Random Forest Classifier works on the Bagging principle, an ensemble technique that starts by choosing a random sample, or subset, from the entire dataset and also a random subset of the input variables. On each subset an individual decision tree is trained, that produces a result. By combining the results of all models through majority voting the most commonly predicted outcome among the models is selected as the final output.

In order to obtain optimal performance from the Random forest, certain conditions must be considered. Firstly, a careful selection of variables should be made in order to avoid the creation of overly correlated decision trees. Furthermore, the number of decision trees to be used depends on the complexity of the problem and the size of the dataset, so that too few trees may lead to an under-robust model, while too many may lead to excessive computational complexity. The depth of the decision trees must also be chosen carefully to avoid overfitting on the training data. Too high a depth can lead to overtraining, while too shallow a depth can lead to an inaccurate model. Finally, the size of the training subset must be chosen so that each decision tree has an adequate number of training examples and sufficient diversity compared to the other decision trees.

3.3. Support Vector Machine

A Support Vector Machine (SVM) is a supervised Machine Learning model, capable of performing linear or nonlinear classification, regression, and even outlier detection. The primary objective of a SVM algorithm is to identify the optimal hyperplane that divides the data points of different classes. SVM tries to find the hyperplane that maximizes the *margin*, which is the distance between the hyperplane and the closest data points of both classes. These closest data points are called *support vectors*. The larger the margin, the better the generalization of the model to new data.

In Figure 1 the two classes can clearly be separated easily with a straight line; they are linearly separable. Sometimes, the data is not linearly separable (i.e. no straight line or plane can separate the classes). In this case, SVM uses a mathematical function called a *kernel* to transform the data into a higher-dimensional space where it becomes separable. If we impose that all instances are clearly separated by the hyperplane, this is called hard margin classification. However, there are two main issues with hard margin classification. First, it only works if the data is linearly separable, and second

it is quite sensitive to outliers. So, to find a balance between maximizing the margin and minimizing classification errors you can use the C parameter that controls this trade-off. A small C allows some misclassifications (soft margin), while a large C prioritizes reducing misclassifications (hard margin).




Figure 1. Margin Maximization Principle

SVMs are potentially designed for binary classification problems. However, with the rise in computationally intensive multiclass problems, several binary classifiers are constructed and combined to formulate SVMs that can implement such multiclass classifications. Moreover, SVM can handle datasets with a large number of features, making it powerful for problems with many variables and they are more robust to overfitting compared to other models. SVMs are particularly well suited for classification of complex but small or medium-sized datasets. For very large datasets, SVMs can be slow to train and may scale poorly, particularly when using non-linear kernels.

3.4. Artificial Neural Networks

Artificial Neural Networks (ANNs) are inspired by the way biological brains process information. In the brain, neurons act as the fundamental units that transmit electrical signals. These neurons are organized in complex networks, with each neuron connected to many others through synapses. They receive inputs, process them, and generate outputs when the accumulated signal surpasses a certain threshold. This behavior is regulated by an *activation function*, which dictates the output based on the weighted sum of inputs. In a similar way, an ANN is a computational model designed to learn patterns from data. It consists of multiple layers of interconnected nodes (analogous to neurons), where each node performs a mathematical operation on the incoming data and passes the result to the next layer in the network. This hierarchical structure allows the network to learn complex representations from raw data, enabling tasks like classification, regression, and pattern recognition. ANNs rely on various functions to process data, and one of the most widely used is the activation function. A common example is the *sigmoid function*, also known as the logistic function, which is expressed as:

$$\sigma(x) = \frac{1}{1 + e^{-x}} \quad (2)$$

This function maps the input to a value between 0 and 1, providing a way to introduce non-linearity into the model. This non-linearity is essential for the network to model complex patterns and make accurate predictions, especially in deeper networks.

When ANNs are extended to include many layers of hidden nodes, they form a deeper architecture known as *Deep Learning*. Deep learning has revolutionized AI by allowing the automatic extraction of complex features from data through multiple levels of abstraction. Deep learning models have been particularly successful in handling unstructured data, such as images, audio, and text, because the additional layers enable them to learn intricate hierarchical features. For instance, in image recognition, early layers might learn to detect edges, intermediate layers might detect textures or shapes, and deeper layers might recognize complex objects like faces or animals.

Training an artificial neural network involves adjusting the model's parameters (i.e., the weights and biases) to minimize the difference between predicted outputs and true outputs (also known as the loss or error). The goal is to find the optimal set of parameters that best fits the data, a process that requires optimization. The most common optimization method used in training neural networks is *Gradient Descent*. Gradient descent is an iterative optimization algorithm that updates the model's parameters in the direction of the negative gradient of the loss function. The gradient represents the rate of change of the loss with respect to each parameter, and moving in the opposite direction helps to reduce the loss. However, basic gradient descent can be slow and inefficient. To address these issues, more advanced optimization algorithms have been developed, including *ADAM* and *SPSA*. In standard gradient descent, the parameters of the model are updated by subtracting a fraction (known as the *learning rate*) of the gradient from the current parameter values:

$$\theta_{new} = \theta_{old} - \eta \cdot \nabla_{\theta} J(\theta) \quad (3)$$

where:

- θ represents the parameters (weights and biases),
- η is the learning rate,
- $\nabla_{\theta} J(\theta)$ is the gradient of the loss function $J(\theta)$.

While effective, basic gradient descent can suffer from problems such as slow convergence and getting stuck in local minima, especially in high-dimensional spaces like those found in deep neural networks.

ADAM is an extension of gradient descent that adapts the learning rate for each parameter individually by considering both the first moment (mean) and second moment (variance) of the gradients. This approach helps accelerate convergence and improve stability, especially when gradients are sparse or noisy. Another interesting variant of gradient descent is SPSA. Unlike traditional gradient-based methods that calculate gradients analytically or numerically, SPSA estimates the gradient through perturbations. This makes SPSA particularly useful in situations where the loss function is expensive to compute or when gradients are noisy. In SPSA, rather than evaluating the loss function at each parameter's current value and its perturbation in each direction, the algorithm performs random perturbations to all parameters simultaneously. This reduces the number of evaluations needed to estimate the gradient, making SPSA more computationally efficient for high-dimensional optimization problems. The general update rule for SPSA is:

$$\theta_{t+1} = \theta_t - \alpha_t \cdot \Delta_t \quad (4)$$

Where:

- Δ_t is the random perturbation at iteration t ,
- α_t is the step size or learning rate.

3.5. Restricted Boltzman Machine

Restricted Boltzmann Machine (RBM) is a type of artificial neural network that is used for unsupervised learning. It is a type of generative model capable of learning the probability distribution

of a given set of input data. This neural network comprises two layers of neurons: a visible layer and a hidden layer. The visible layer corresponds to the input data, while the hidden layer captures the features learned by the network. The term "restricted" refers to the fact that neurons within the same layer do not connect to each other. Instead, each neuron in the visible layer is linked only to neurons in the hidden layer, and vice versa. This structure enables the RBM to create a compressed representation of the input data by reducing its dimensionality.

4. Quantum Machine Learning

ML is inherently filled with complexity and large size computational tasks, so that the extraordinary properties of quantum computing have been naturally associated with the ML domain. Hence, Quantum Machine Learning (QML) is one of the most fascinating and promising applications of quantum computing. The concept of QML encompasses a broad variety of methods and algorithms that are still under investigation. Basically, machine learning routines could benefit from quantum computing in two distinct ways: (i) *quantum data encoding*: systems of qubits constitute the quantum Hilbert space and in principle QML could solve learning tasks much faster by accessing this exponentially large dimension space; (ii) *quantum data processing*: theoretical computational speedups of several quantum algorithms have been proven and many of them could have high impact in solving certain ML tasks such as sampling, linear algebra, and Fourier transforms [25,52].

Despite the immense potential, there are some limitations of using QML in practical applications, mostly due to the level of maturity of the hardware [104]. The first big caveat is the existence of a Quantum Random Access Memory (QRAM), which is still a debatable topic of research. Many quantum algorithms require a "reading cost" when accessing data and this could even be dominating in the end-to-end computation, neutralizing the speedup during processing. Indeed, QRAM has been already theoretically designed [44,141], and its concrete implementation would enable huge advancements in accessing data directly from quantum hardware. The second challenge is the actual noise when performing quantum computation tasks. The current quantum hardware is characterized by low qubit availability, scarce coherence times and inherent impurities, which leads to noisy results of the computation. This is why one of the main fields of research is how to implement an efficient quantum error correction (QEC) [125]. Most of the best quantum speedup algorithms require severe error thresholds plus extensive hardware capacity, and this is the first reason that prevents them from being used in real applications.

Indeed, even with the current hardware, the so-called Noisy Intermediate-Scale Quantum (NISQ) devices, many near-term quantum machine learning algorithms have been developed. Most of these methods account for a variational approach, that is a hybrid quantum-classical workflow where the quantum processing unit is used only for a specific subroutine, the one that can benefit the most from quantum computing. On the other hand, the rest of the computation, held by classical and/or high-performance computing, is meant to optimize a set of parameters (this is why it is called variational approach) so as to minimize the cost function whose minimum leads to the solution of the problem to be solved. In this respect, one can compare variational quantum algorithms to state-of-the-art Deep Learning models, where a sort of quantum network is optimized via complex parametric tuning, and trained to solve various tasks. Ultimately, there is also another way to quantum machine learning: using quantum-inspired algorithms that employ classical data and classical hardware could still give a certain benefit in the immediate future as a tactical solution [21].

Since the ultimate goal of this review is to give a view on near-term and state-of-the-art applications of quantum machine learning, the following sections are mainly focused on variational quantum algorithms and other "early-available" quantum methods applied to machine learning and their applications to the energy sector.

4.1. Quantum Variational Algorithms

Variational quantum algorithms are based on two essential ingredients: (i) the definition of a *cost function* (or loss function) that measures how well the algorithm is doing, and (ii) *ansatz circuit*, which

is a set of quantum operations (in a gate-based form) that describes a good solution to the problem, with its tunable parameters adjusted accordingly. Without loss of generality, a parametrized quantum circuit (PQC) ansatz can be expressed by a unitary transformation $U(\theta)$, where $\theta \in \mathbb{R}^p$ is the vector of the parameters. Essentially, a variational quantum algorithm (VQA) wants to minimise a function $C(\theta) : \mathbb{R}^p \rightarrow \mathbb{R}$ as such: $\theta_{opt} = \arg \min_{\theta} C(\theta)$. Typically, the “quantum-form” of this cost function can be expressed as:

$$C(\theta) = \langle O \rangle_{\theta} = \text{Tr}[OU(\theta)|0\rangle\langle 0|U(\theta)^{\dagger}] \quad (5)$$

that is a function of the parameters, the specific ansatz, and the observable. There are some important guidelines to determine a good choice for a loss function: (i) classically hard to compute but accessible to quantum devices: this is the most important point when seeking quantum advantage; (ii) trainable: there should exist an iterative method capable of finding the global minimum of this function with a certain efficiency; this property is essential in order to guarantee the efficacy of the whole variational procedure.

One fundamental step which is strictly referred to quantum computing principles is the choice of the ansatz $U(\theta)$. Following the variational procedure, this choice has to be done “a priori”, respecting a nice balance between complexity (to express optimal solution) and practicality (to favour convergence). Generally, most of the ansatzes have a layered-factorial form

$$U(\theta) = \prod_{\ell=1}^L W_{\ell} U_{\ell}(\theta_{\ell}) = W_L U_L(\theta_L) \dots W_1 U_1(\theta_1), \quad (6)$$

where W_{ℓ} is the fixed encoding part and $U_{\ell}(\theta_{\ell})$ is the actual variational ansatz.

Among the several existing strategies to compute an effective ansatz, there are: (i) *hardware-efficient ansatz* (HWE), that is meant to be easily run on current, near-term quantum hardware [72]; (ii) *problem-inspired ansatz*, that is built by leveraging the knowledge of the underlying physics (when present) [89]; (iii) *adaptive ansatz*, which is specifically designed to adaptively optimize its structure along with that of the parameters [139]. As far as the classical part of a cost function is concerned, optimal choices of both classical observables and optimization methods have to be done. The first is strictly related to the ultimate objective of the problem (i.e. classification, regression, etc.), so that the observable O is chosen in accordance to the nature of the problem. Secondly, one has to choose among the different options on iterative optimization methods. One effective choice for such differentiable functions is to use gradient-descent loop, which is already widely studied in classical machine learning literature [13].

4.2. Expressivity-Trainability Trade-Off

The most impending challenge for VQAs is the one of the “barren-plateau”: it is rigorously shown that, under mild hypothesis, the average gradient of a random PQCs is equal to zero [20,90]. Not only this, but the gradient variance of a random PQCs decays exponentially with the number of qubits [56]. This problem is also called “vanishing gradient” and could be really detrimental especially in the regime where qubits are much more than 1, which is, by contrast, when one would expect the most quantum advantage. Fortunately, there are many mitigation strategies [40,84,100]. All of them are basically focused on moving away from the vanishing gradient initial hypothesis of a “universal”, random unitary that leads to maximum expressivity but difficult trainability at the same time. The idea is to have an ansatz where complexity is controlled to address specific, problem-inspired expressivity. Thus, trainability is preserved and efficiency optimized.

Another way to improve trainability compensating for the vanishing-gradient phenomenon is to work on classical optimization methods: one can either operate on optimizers like ADAM or SPSSA algorithms or employ a varying learning rate over training. Furthermore, since the appearance of the barren-plateau is strictly related to the variational parameters initialization [73], techniques like the so-called “warm-start” have been successfully employed to adjust parameters initialization in

a pre-training phase [110]; similarly, one strategy is to optimize a subset of parameters at a time in a layer-wise approach [128]. Another completely different approach is to work without a call to a classical optimization algorithm: for any variational cost function that is a linear sum of expectation values, an exact formula can be extrapolated to find one optimal parameter at a time that guarantees a decreasing behavior of the cost function over training [71].

In summary, the problem of barren-plateau is a critical topic for near-term quantum ML applications and although there is no universal solution available, there are many mitigation strategies that depend on the specific case that can be adopted.

4.3. Explicit Quantum Models

Suppose to have a set of classical data $\{x_i\}_i^m = 1$. The first fundamental step to VQA implementation is the so-called "feature-embedding", where these data are mapped to the quantum Hilbert space. This is done by using a unitary operation $U(x)$, that maps input data to the quantum Hilbert space:

$$\phi: \mathcal{X} \rightarrow \mathcal{F}, \quad \mathcal{X} \subset \mathbb{R}^d, \quad \mathcal{F} = \mathcal{H} = \mathbb{C}^{2^n} \quad (7)$$

$$x \mapsto |\phi(x)\rangle = U(x)|0\rangle^{\otimes n} \quad \text{or} \quad x \mapsto \rho(x) = |\phi(x)\rangle\langle\phi(x)| \quad (8)$$

$\mathcal{F} = \mathcal{H}$ being the qubits Hilbert space of the quantum computer. The operation $U(x)$ can be chosen among different alternatives [85]. The most frequently used are: (i) *amplitude embedding*: encodes a normalized N-dimensional feature vector into a quantum state of n-qubits. This is done by putting the normalized data as amplitudes of a quantum wavefunction:

$$|\psi_x\rangle = \sum_{i=1}^N x_i |i\rangle \quad (9)$$

where $N = 2^n$, x_i is the i -th element of x , and $|i\rangle$ is the i -th computational basis state; (ii) *angle embedding*: encodes a n-dimensional sample into an angle rotations applied to as many qubits

$$U(x) = \bigotimes_{i=1}^n R_{P_i}(x_i), \quad x = [x_1, x_2, \dots, x_n], \quad P_i \in \{X, Y, Z\} \quad (10)$$

Single gate operations as Pauli rotations $R_P(x)$ are Pauli rotations (2.6) around one of the Pauli axis $P \in \{X, Y, Z\}$ are parametrized by the input values and then used to map the resulting quantum state; (iii) *basis embedding*: encodes a N-bit binary string input $x^{(m)} = (b_1, \dots, b_n)$ with $b_i \in \{0, 1\}$ into a quantum state of N qubits. This is achieved by simply mapping the bits to a corresponding computational basis state, with the formula:

$$|D\rangle = \frac{1}{\sqrt{M}} \sum_{m=1}^M |x^{(m)}\rangle \quad (11)$$

From these three different approaches, many variants have been developed [63,71,86]. Balancing robustness with as few qubits as possible is crucial to address real problems with the near-term hardware availability. After the encoding phase, the variational unitary $V(\theta)$ is applied. The resulting quantum state follows:

$$|\phi_\theta(x)\rangle = V(\theta)|\phi(x)\rangle = V(\theta)U(x)|0\rangle^{\otimes n} \quad (12)$$

Then, the last step of an explicit quantum model is to compute an observable at the end of the quantum computation:

$$f_\theta(x) = \langle O \rangle_{\theta;x} = \langle \phi_\theta(x) | O | \phi_\theta(x) \rangle = \langle \phi(x) | \underbrace{V(\theta)^\dagger O V(\theta)}_{O_\theta} | \phi(x) \rangle \quad (13)$$

This observable is directly used in the loss function of a supervised machine learning model, i.e. a classification or a regression problem [19,53], to drive the training procedure. Notably, the optimization

of the variational parameters is implemented by directly measuring the observable on a quantum computer.

4.4. Implicit Quantum Models (Quantum Kernels)

Similarly to classical Kernels (discussed in sec. 3.1), Quantum Kernels can be defined [94]. These methods (often called *implicit models*) work as the classical framework where now the Kernel function is given by the inner product in the quantum feature space:

$$K(x, x') = |\langle \phi(x) | \phi(x') \rangle|^2 \quad (14)$$

Thus, quantum computers can be only used to compute the Kernel function, whereas the variational parameters that define the optimal model as a linear combination of kernel evaluations are purely classical. This quantum kernel estimation (QKE) has been first proposed by Maria Schuld and Nathan Killoran[117]. Shortly after that, a subsequent work proved that there exist classically intractable kernels that can be accessed with the use of the quantum feature space with a theoretical exponential advantage [53]. From that on, several quantum circuits have been proposed to calculate quantum kernels applied to machine learning models. Particularly, quantum kernels can be used within regression or classification tasks as, for example, quantum support vector machines [42,74]. Notably, as for their classical version, implicit models are guaranteed with the best performance by the *representer theorem* [116]. On the other hand, they require $O(m^2)$ queries to the quantum computer where m is the number of samples in the training dataset, plus an additional $O(m^3)$ classical-processing queries to compute the optimal Kernel weights for the optimal predictor [70]. By contrast, explicit models need $O(pm)$ calls to the quantum computer, where p is the number of variational parameters that have to be tuned. Evidently, when there is no need to have a large number of parameters to describe the variational block $V(\theta)$, explicit models are cheaper and therefore more efficient than implicit ones. By contrast, quantum kernels are proven to guarantee advantage under specific, complex conditions.

4.5. Quantum Neural Networks

Explicit and implicit models have in common the fact that they can be mathematically treated as an inner product between a data-dependent quantum state (or feature map) and a trainable observable [122]. This property gives the previous methods a linearity framework to lay on. When this linearity is broken, similarly to what happens in classical machine learning, we enter the realm of quantum neural networks (QNN). Differently from the two previous approaches, a quantum neural network is defined when a parametrized quantum circuit uses a repeated structure of both encoding and variational layer, assuming the form:

$$U_\theta(x) = V_L(\theta_L)U_L(x) \dots V_1(\theta_1)U_1(x)V_0(\theta_0) = \prod_{l=1}^L U(x)V_l(\theta_l) \quad (15)$$

where L is the number of layers of the QNN. The fundamental property of this form of unitary is to add "*data re-uploading*", that is the repeated use of input data in the unitary computation via multiple encoding layers. Notably, although the name recalls the classical neural network framework for the layered structure, QNNs are very different in terms of approach and underlying mathematics. Theoretically, this type of PQC is proven to be a different form a *L-truncated Fourier series* when angle-embedding encoding is employed [43]:

$$f_\theta(x) = \langle O \rangle_{x;\theta} = \langle 0 | U_\theta(x)^\dagger O U_\theta(x) | 0 \rangle = \sum_{\omega \in \Omega} c_\omega e^{-i\omega \cdot x} \quad (16)$$

where $\Omega = \{\omega \in \mathbb{R}^d\}$. Such formalism is very helpful because it gives a way to characterize QNN generalization performance by the use of already known statistical concepts. Essentially, it has been proven that the Generalization Error (GE) of a QNN can be bounded as [18,118]:

$$GE \leq \mathcal{O}\left(\sqrt{\frac{|\Omega|}{m}}\right) \quad (17)$$

where Ω is the Fourier spectrum of frequencies accessed by the QNN and m is the training data size. Notably, when Ω is too large (which means a high complex quantum circuit), the model weakens in such a measure that only a corresponding increase of data size is able to compensate. This consequence is something one has to take into account when designing a QNN. Given that, several studies of practical applications have shown a practical advantage in terms of expressivity of QNN when compared to classical versions of the same size, raising interest on the topic [45,63,101]. Of particular importance is the fact that QNN structure can be integrated in classical NN architectures, forming hybrid quantum-classical neural networks. This approach allows a greater scalability when large datasets are employed, that is basically the case for almost every real-world scenario. For instance, deep learning and VQA can be trained and run in parallel where classical NN layers can be used to further reduce the input data dimension so that VQA's employment is enabled with the current quantum hardware available. Although there can be no analytical proof of the effectiveness of the quantum integration in such hybrid structures, applications of hybrid QML are widespread in literature with evidence of empirical advantages over purely classical counterparts [12,81,119].

4.6. Quantum Annealing Applied to Machine Learning

Indeed, apart from the variational approach employing gate-based hardware, there are alternative “quantum” ways to enhance classical machine learning performance, one of them being *quantum annealing*. Since Machine Learning is filled with optimization problems, quantum annealing processors have become popular in solving some of these optimization tasks in machine learning applications. One of the first and most popular fields of application for QA is feature selection, that is the process of extracting the most efficient subset of features from input data of ML models. This is a complex decisional problem that can be solved with Ising-like models run on annealers. Quantum hardware is therefore potentially helpful when dealing with wide and complex datasets [37]. One other interesting use of QA is the optimal sampling of data to train random forest (RF) classification models [132], which has been proven able to increase classification accuracy of the purely classical RF. Furthermore, a popular quantum-enhanced deep learning model is the restricted Boltzmann machine (RBM). Quantizing the classical version of the RBM is simple: one has to take the network and express it as a set of coupled spins, resembling the Ising model. Then, the last one can be solved using QA, obtaining the optimal, tuned network. Quantum RBM models have been successfully implemented in a real-scenario classification problem with interesting results [30,96].

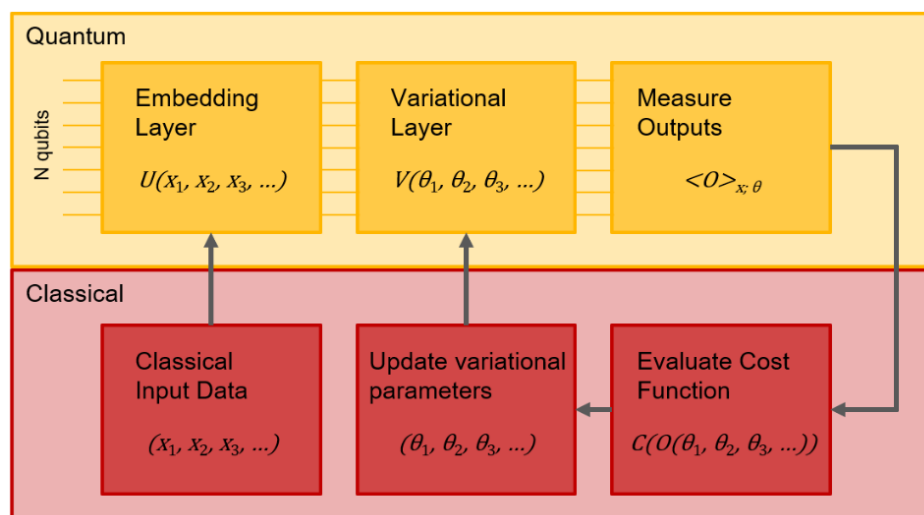


Figure 2. Representation of the high-level workflow for a quantum variational algorithm.

5. Case Studies in E&U

Electricity is arguably the most important energy resource of modern society. Among the essential daily activities that depend on electricity one can name healthcare services, communications, information technology, house devices usage, and many more. All this is guaranteed thanks to a reliable and continuous supply from power plants. The major stages of a power plant process are: generation, transmission, distribution. Today electricity operators companies are facing old and new challenges to keep up with society's needs. On one hand, the global energy demand is destined to increase, requiring the extension of existing power supply systems. On the other hand, environmental impact must be taken seriously into account, shifting to more sustainable processes. One (if not the only) viable solution is the adoption of key-enabling new technologies, capable of dealing with today's increasing complexity but rooted on sustainability as a paradigm. Among these, one of the most important is undoubtedly artificial intelligence, whose efficacy is determined by machine learning algorithms and high-performance computation. That is why SPOKE 10 of NextGenEU program of funds, managed by Politecnico di Milano, is dedicated to high performance and quantum computing (HPC and QC) [24]. QC is particularly interesting for its disruptive potential in the energy and utilities (E&U) industry, which is filled with computational complexity. In particular, forecasting problems are crucial to establish efficient power supply operations and they mostly rely on ML. Fortunately, as exposed in previous sections, QC in its theoretical framework suits nicely with ML. Although the spectrum of quantum machine learning applications for E&U is even broader than the following review, the focus of this work is on quantum applications that can be practically assessed with today's technology maturity. The reason for this pragmatic approach is two fold: first, the objective is to shed light on early opportunities for quantum technology adoption in a reasonable horizon of time; secondly, only when real hardware faces real problems the actual capabilities of a new technology are really explored. With this in mind, quantum machine learning near-term applications are briefly reviewed below.

5.1. Method and Overview

Given the growing interest in quantum technologies and their huge potential for solving complex problems, the energy industry is undoubtedly one of the principal candidates to benefit most from quantum computing advancements. Problems related to energy are inherently very complex and their impact on society makes it urgent to address the strategic objectives for the industry, such as meeting growing demand and reaching sustainability. For many years machine learning and AI techniques have been employed successfully to solve some of the most important problems of forecasting, classification and imputation. Even the energy industry is currently benefiting from the rapid growth of ML and AI

in their value chain. However, in many applications, having a faster or more accurate solution could be strategic. This is why quantum computing applied to machine learning is a promising frontier to explore for the energy industry. With these premises, we formulated two research questions:

- How can the energy and utilities sector benefit from quantum computing, and which specific ML applications or challenges will QC address in the near- to medium-term future?
- Which use cases of quantum machine learning have the most significant impact on the energy and utilities sector related to their level of readiness?

To answer these questions, we gather a collection of key studies with specific inclusion and exclusion criteria. Peer-reviewed papers with the date of publication before February 2025 have been included in the analysis. As inclusion criteria, there have been the following: (i) the application involves one or more computational problems of the energy industry value chain; (ii) the algorithm implies a machine learning technique; (iii) the quantum computing solution proposed is viable in the NISQ era of quantum computers, which include gate-based variational approach, hybrid architectures and quantum annealing. We have also implemented the following exclusion criteria: (i) employment of general-purpose quantum algorithms that imply scaled fault-tolerance quantum computers (i.e. Grover algorithm, HHL, etc.); (ii) the absence of a real test on datasets referred to the energy industry or related applications; (iii) quantum-inspired approaches.

Key studies that have been reviewed in the following clearly suggest that quantum machine learning has the potential to be applied to a wide range of applications in the energy industry. The number of total articles reviewed is 23, originating from 32 different countries (each country associated with an author's institution is considered individually). The main source of articles is the USA, followed by China. By continent, Asia is overall more present. Four main categories of applications have been identified which belong to a piece of the energy industry value chain: electrical load management, energy generation forecasting, grid operations and maintenance, and finance for sustainable energy. A set of descriptive features of these works has been extracted: (i) *method*, that is the specific name that authors gave to the quantum model they use; (ii) *typology*, which is a closed choice between the five categories of quantum computing methods reviewed in chapter 4: 1. Explicit VQA; 2. Implicit VQA (or Quantum Kernel); 3. Data re-uploading (or QNN); 4. Hybrid (mix of PQC and classical workloads in a single model architecture); 5. Annealing (or analog quantum computing); (iii) *software technology*, that is the reference (when present) of which cloud platform/SDK/libraries have been used for the application; (iv) *hardware technology*, that is the reference of which QPU has been used (we indicate the simulator in case of no QPU employed).

As one can see from the "Look Up Table for Use Cases" (Table 1 in Appendix), the category with the highest number of studies is related to energy generation forecasting, which suggests that this is one of the hottest topics for the industry. Additionally, data collection for those types of use cases is made easy by the availability of transparency data on energy generation from plants and weather data. Nonetheless, given that many different methods of generation forecasting are employed for the same application, the major category of distinct use cases is the one related to grid operations and maintenance. This category is indeed very rich in applications and, additionally, use cases on the energy transmission value chain are often the most impactful for energy industry businesses, with major consequences for society and final users.

Table 1. Look Up Table for Use Cases

Value Chain	Category	Use Case	Reference	ID
Distribution	Electrical Load Management	Load Forecasting for Demand Response	[97], [112]	1
		Automated Demand response in Smart Cities	[4]	2
		HVAC Automated Control in Buildings	[9]	3
Generation	Energy Generation Forecasting	Solar Irradiation Forecasting	[120], [140], [98], [60], [130]	4
		Photovoltaic Power Forecasting	[113], [75]	5
		Wind Speed Forecasting	[59]	6
		Forecasting Power from Off Shore Wind Farms	[50]	7
		Climate Time Series Forecasting	[61]	8
Transmission	Grid Operations and Maintenance	Fault Diagnosis in Electrical Power Systems	[3]	9
		Photovoltaic Panels Fault Detection	[135]	10
		Wind Turbine Pitch Fault Detection	[26]	11
		Power System Stability Assessment	[143]	12
		Dynamic Event Identification Using Phasor Measurement Units in Power Systems	[68]	13
		Smart Grid Stability Forecasting	[51]	14
		Defect Detection in Wind Turbine Gearbox	[41]	15
Financial Operations	Finance for Sustainable Energy	Carbon Price Forecasting	[17]	16
		Optimal Scheduling of EV Recharges	[9]	17
		Carbon Market Risk Estimation	[142]	18

5.2. Electrical Load Management

5.2.1. Overview

Predicting energy demand with reasonable accuracy is arguably one of the oldest and most important challenge for a power system supply chain. Modeling future information on energy demand is a rich and wide field of research. First of all, energy demand forecasting methods can be categorized by two essential features: spatial and temporal resolution. The former describes models ranging from single appliance to country-level energy demand, while the latter varies from sub-hourly energy request to yearly energy consumption. More local and responsive forecasting is crucial to daily demand-response operations in power grids and day-ahead energy markets. In turn, more aggregated and long-term information are useful for strategic planning and decision making. There are several specialized reviews that cover the range of existing techniques for energy demand estimation [7,28,58,80,87], which we refer to for an extensive view on the topic. Here we mention the most important criteria of an energy forecasting method: spatial resolution, temporal resolution, input features, data preparation, data size, applied techniques, forecasting horizon, metrics of performance. In the following the focus will be on the range of the applied techniques, starting from classical and conventional methods and eventually reaching the use of quantum computing.

The simplest classification of load forecasting techniques in literature is to divide them between conventional techniques and machine-learning-based techniques [137]. The former class involves statistical methods, subdivided in time series analysis (TSA) and regression methods. TSA techniques consist in using the historical time series of data to infer future predictions. Popular TSA models are autoregressive moving average analysis ARMA and ARIMA, that is its version for non-stationary time series [14]. On the other hand, regression models are meant to find a functional relation between inputs and outputs of a system, by minimizing a cost function that is typically a figure of optimal fitting between data points and the regression wave of predictions. In the case of the linear relation hypothesis, the cost function is the sum of the squares of the residuals, that are the difference between observed values and the fitted ones [121]. Then, non-linear functional relations can be assumed with the use of more complex cost functions or implementations of Kernel methods [27]. An advancement from regression methods are machine learning models. NN and SVM are the most employed ML models to derive the function that relates input data to the energy load demand historic output. Once this relation is trained, learned and validated, current input data are fed to the model so that future energy demand can be efficiently predicted. In the category of ML-based techniques one also has decision trees (DT) [93], Bayesian [133] and ensemble learning [2] approaches. Other than conventional and ML-based techniques, there are also other alternative methods applied to energy demand forecasting. One of them is the employment of metaheuristics combined to ML: genetic algorithms, particle swarm optimization and other similar algorithms can be used to optimize parameters search in SVM or NN models [134]. Furthermore, stochastic processes can be employed to simulate load profile, using Markov chains [8], while more advanced techniques based on probability theory like fuzzy logic and grey systems are a good way to develop uncertainty aware forecasting models [62,67].

5.2.2. Key Studies

QML is an interesting perspective for more accurate and resolute energy demand forecasting. In their work, Kumar et al. discuss the application of Quantum Support Vector Machine (QSVM) in forecasting energy consumption within Home Energy Management Systems (HEMS) in smart grids [97]. The authors highlight the importance of accurate load forecasting to manage energy consumption efficiently, which is crucial due to the fluctuating energy demands in households. The study compares QSVM with traditional deep learning models like Recurrent Neural Networks (RNN) and Long Short-Term Memory (LSTM) networks, demonstrating QSVM's superior performance in handling complex and nonlinear electricity consumption patterns. Their methodology involves using the AMPds2 dataset, which contains data on various household energy loads. After data pre-processing and normalization, a quantum kernel is computed using 2, 4 and 8 qubits. Angle embedding and ZZ

feature map [131] as circuit ansatz are employed to model the kernel function. Once computed, the quantum kernel is used for an SVM conventional training. The results on testing data show that QSVM achieves a higher accuracy of 97.36% compared to the deep learning models, which have slightly lower accuracies. This is attributed to QSVM's ability to efficiently process large datasets and capture intricate relationships within the data. This outcome emphasizes QSVM's potential as a promising solution for enhancing forecasting accuracy in HEMS, suggesting that future research could explore its adaptability to different datasets and its performance in real-world scenarios. The study also proposes the development of hybrid models that combine multiple forecasting techniques to further improve prediction accuracy.

More recently, Safari and Badamchizadeh introduced NeuroQuMan, a cutting-edge approach that combines quantum computing and machine learning to enhance energy management systems [112]. This innovative method employs a three-Qubit quantum neural network (QNN) to predict energy demand based on user reaction times, aiming to optimize energy consumption and improve grid stability. The structure of NeuroQuMan involves several key components: data preprocessing, quantum device initialization, quantum circuit definition, user decision-making, QNN predictions, and visualization of results. The process begins with preparing real-time data, such as reaction times and power outputs, for analysis. A quantum device is then initialized to execute quantum circuits, which are defined using specific quantum operations. These circuits simulate user decision-making and optimize power generation through economic dispatch calculations. In detail, the architecture followed the typical structure of a variational QNN: firstly, a quantum data encoding of response times, energy loads and power generated was performed via 3-qubit angle embedding. The actual scheme of a variational quantum circuit is then iteratively applied using quantum gates. Finally, the output is measured from the final quantum state. The QNN model predicts demand and renewable energy generation, with the results visualized to assess performance. The results from simulations show that NeuroQuMan achieves significant accuracy in predicting demand response, with metrics like RMSPE and MAPE indicating its effectiveness. The model's performance is compared with other methodologies, demonstrating its superior accuracy and efficiency. Despite its promising results, NeuroQuMan faces challenges such as the limitations of current quantum hardware and the complexity of quantum circuit design. Future work aims to integrate advanced quantum machine learning techniques, like Quantum LSTM and Quantum Neuro-Fuzzy systems, to further enhance prediction accuracy and manage uncertainties in energy systems.

Concerning hybrid architectures, a proper work from Ajagekar and You [4] discusses a cutting-edge approach to managing energy consumption in buildings using a hybrid quantum-classical control framework. This innovative strategy leverages the strengths of both quantum and classical computing to enhance demand response, which is crucial for reducing energy use and carbon emissions in grid-interactive buildings. The study's primary objective is to develop a robust control system that can efficiently manage energy demands in buildings equipped with renewable energy sources and storage devices. By using variational quantum circuits (VQCs) within a reinforcement learning framework, the system can make real-time decisions to optimize energy use. This approach addresses the limitations of traditional methods, offering improved scalability and performance as the size of building microgrids increases. This controller is formulated as a markov decision process and is based on three main components: (i) estimation of the value function using VQC: the VQC is used to approximate the value of the state-action function for a state and control pair (x, u) . In a reinforcement learning context, the value function represents the cumulative sum of discounted rewards over an infinite horizon; (ii) extraction of controls from a trained VQC: after the VQC has been trained, it is used to extract the necessary controls to manage energy storage devices and optimize the energy consumption of buildings; (iii) integration of quantum and classical routines: this enables circuit learning and the calculation of optimal control. The VQC consists of three main components: first, an amplitude encoding scheme responsible for translating a classical data into a quantum state. Then, parametric gate operations are applied to the input quantum state to manipulate it according to the trainable

parameters of the VQC. These are obtained by combining an entangling layer composed by CNOT gates, single qubit Pauli rotations, and lastly the pooling operator is applied, comprising a sequence of operations between a source and sink qubit pair, like sink qubit inversion along z-axis, inverted CNOT among the two, plus fixed amount rotations along z- and y-axis. Ultimately, the circuit learning phase measures the resulting quantum state with an observable to estimate the objective function. Learning the function is formulated as minimizing a loss function, and a gradient descent algorithm is used to update the parameters of the variational circuit. The results of the study are promising, showing that the hybrid VQC-based strategy can reduce energy consumption and carbon emissions by over 13.6% compared to classical methods like deep deterministic policy gradient and model predictive control. The framework is designed to be adaptable, handling uncertainties such as fluctuating renewable energy generation and varying carbon intensity levels. The hybrid approach is proven to be a valid way to enhance computational efficiency and address scalable solutions for current quantum hardware. This work underscores the importance of quantum-classical integration of architectures, combining quantum workloads with classical optimization to tackle complex and critical real-world challenges like adaptive response to energy demand.

Another work on quantum reinforcement learning is proposed by Andrés et al. [9]. The document outlines a methodology for training RL agents using a hybrid classical and quantum approach. The primary focus is on the design and training of agents in an environment called Eplus-demo-v1, utilizing the Advantage Actor-Critic (A2C) algorithm. Eplus-demo-v1 simulates a heating, ventilation, air-conditioning HVAC scenario in a target building, while the RL algorithm's main objective is saving and optimizing energy demand. The methodology involves two distinct types of experiments: one using a classic feedforward multilayer perceptron (MLP) neural network and the other employing a hybrid architecture that combines a linear layer with a variational quantum circuit (VQC). In the classical approach, the MLP neural network is used to model the agent's policy. This network is trained using the A2C algorithm, which is a popular method in RL for learning stochastic policies. The training process involves running four environments in parallel to generate batches of interactions between the agent and the environment. A discount factor of 0.98 is used to calculate the total return of each episode, and the training is set to stop after 100 episodes. The agent's performance is evaluated in the final episode using a deterministic policy that selects the action with the highest probability. The quantum approach involves designing a quantum agent with a hybrid architecture. This agent also uses the A2C algorithm for training, maintaining the same settings as the classical agent to ensure a fair comparison. The quantum agent's policy is implemented with a variational quantum circuit, which is integrated with a linear layer. This configuration aims to leverage the potential advantages of quantum computing, such as reduced parameter space and computational complexity, while still utilizing classical neural network components. The authors also discuss the structure of the networks used in the experiments. For the classical MLP networks, the number of parameters is significantly larger compared to the quantum VQC networks. The quantum networks have a reduced number of parameters, which is expected to decrease the computational complexity of the simulations. The methodology emphasizes the use of separate networks for different action sets, which helps manage the computational demands of simulating quantum neural networks. The results highlight the total accumulated rewards obtained by both quantum and classical agents. The VQC agent consistently converges to the optimal policy, achieving the same value for average, best, and worst total rewards, as well as in test episodes using a deterministic policy. This consistency indicates that the VQC agent effectively learns the optimal policy. In contrast, the MLP agent does not reach the optimal policy, with its best execution yielding worse total and average rewards across all executions. However, the authors also note a limitation of the quantum agent, which is the high computational time required to simulate the quantum environment.

5.3. Energy Generation Forecasting

5.3.1. Overview

Energy Generation Forecasting (EGF) has become increasingly critical in the context of modern energy systems. With the growing integration of renewable energy sources, such as wind, solar, and hydropower, the need for accurate and reliable forecasts is more important than ever. These energy sources, while sustainable, are often intermittent and volatile, requiring advanced forecasting models to predict their output and ensure grid stability, in order to prevent overloads and guarantee a smooth transition to a more sustainable energy ecology. To face this problem several studies for solar and wind power forecasting have implemented both direct and indirect forecasting approaches. The conventional direct method for predicting solar and wind power relies on time-series analysis without considering specific predictions on weather data. This method benefits from not needing models that link that link weather and power data, but its accuracy can be limited due to the inherent randomness of power data contingencies. Another direct method uses meteorological data from Numerical Weather Prediction (NWP) models to forecast solar and wind power directly. On the other hand, several studies have proposed an indirect approach to enhance forecast accuracy that consists in forecasting the expected solar or wind profiles at a specific site and timeframe, and then converting these profiles into power data using appropriate weather-to-power performance models. For all of these reasons EGF plays an important role in effective energy management, balancing supply and demand, optimizing resource usage, and enhancing the resilience of power systems. However, as the energy landscape continues to evolve, especially with the rise of smart grids and decentralized energy systems, the complexity of these forecasts has increased. In the following, we will explore the methods and techniques used in energy generation forecasting, highlighting the challenges involved and examining the latest advancements that aim to improve forecast accuracy and reliability.

As in the previous section, the most commonly used machine learning methods for energy generation forecasting are Artificial Neural Networks (ANNs) and Support Vector Machines (SVMs). Studies highlight that Neural Networks (NNs) outperform traditional time-series models due to their nonlinear fitting capabilities. Over time, NN architectures have evolved from simple models to deep learning (DL) configurations. The main differences between Energy Load Forecasting (ELF) and EGF regard the input of the models used for the forecasts. For example, the researchers have identified that PV power generation mainly depends on the amount of solar irradiance. In addition, other weather parameters, including atmospheric temperature, module temperature, wind speed and direction, and humidity, are considered as potential inputs. Other parameters include air temperature, air humidity, atmospheric pressure, as well as seasonal parameters and historical turbine velocity for wind power forecasting. As for ELF, hyperparameter selection is critical for ML models, as incorrect values can cause overfitting or underfitting. To improve accuracy, researchers have optimized SVM parameters using enhanced Ant Colony Optimization (ACO), genetic algorithms, and swarm-based techniques. Similarly, DL-based PV forecasting models have benefited from Particle Swarm Optimization (PSO) and other advanced tuning methods [69].

Long Short-Term Memory (LSTM) networks are particularly effective in time-series forecasting as they capture long-term dependencies and mitigate gradient descent issues. Studies comparing NN models with traditional forecasting methods found that NNs offer superior predictive accuracy with lower computational costs [123].

Additionally, hybrid models combining ML and physical techniques have been explored for irradiance-to-power conversion. A large-scale study comparing physical, data-driven, and hybrid methods for day-ahead PV power forecasting proposed optimization strategies to minimize forecasting errors (MAE and RMSE), providing a comprehensive approach to improving prediction accuracy [88].

A study using PV infrastructure data from Cocoa, Florida, found that Artificial Neural Networks outperformed other algorithms, achieving the best forecasting metrics (MAE: 0.4693, RMSE: 0.8816 W, R^2 : 0.9988), proving ANN to be the most reliable method for PV power prediction. To further address forecasting limitations, researchers proposed a hybrid framework combining Attention-based

LSTM (A-LSTM) for nonlinear time-series analysis, Convolutional Neural Networks (CNNs) for local correlations, and an autoregression model for linear patterns. This hybrid model significantly improved accuracy over ANNs and decision trees (DTs), reducing MAE by 13.4% for solar PV, 22.9% for solar thermal, and 27.1% for wind power. Additionally, Gaussian Processes (GPs) were explored for probabilistic forecasting, with a new Log-normal Process (LP) model introduced for positive data, such as home load forecasts, addressing a previously overlooked issue [33].

5.3.2. Key Studies

EGF of REs has been immediately seen by researchers as an early application of QML. At the very first stages of quantum machine learning demonstrations, Senekane, Makhamisa and Tael, Benedict Molibeli et al. developed an explicit model using VQA to forecast solar irradiation [120]. The author discusses one of the very first practical use of a quantum support vector machine (QSVM), with a full explicit QML approach composed by state preparation, matrix inversion, angle embedding and a variational classifier [82]. The implementation involved using a dataset from the Digital Technology Group Weather Station at Cambridge University, which included measurements of temperature, humidity, and wind speed. The QSVM was simulated using Python, specifically employing the Scikit-learn library for machine learning tasks and Orange data mining software for visualization. The results showed that the QSVM could effectively predict solar irradiance, with the best performance achieved when 70% of the dataset was used for training.

A more recent work on indirect solar power generation forecasting discusses the application of Quantum Neural Networks (QNN) for predicting solar irradiance using the Folsom dataset from California [98]. The study compares the QNN model with traditional machine learning models like Support Vector Regression (SVR), Group Method of Data Handling (GMDH), and Extreme Gradient Boost (XGBoost). Two feature mapping strategies of the QNN structure were explored: angle encoding and ZZ feature map, with angle encoding using Pauli Y gates proving more effective. The ansatz, which acts like the training weights in classical neural networks, was tested using different configurations, including the Two Local and Quantum Approximate Optimization Algorithm (QAOA) approaches. The Two Local ansatz with linear entanglement and one repetition, optimized using the L-BFGS-B algorithm, provided the best results. The study found that while classical models like XGBoost performed better for shorter forecasting horizons, the QNN model showed promise for longer prediction windows, indicating its potential to process complex spatiotemporal information more effectively.

A novel approach to predicting solar irradiance using a hybrid model that combines Quantum Long Short-Term Memory (QLSTM) networks with Variational Quantum Circuits (VQC) is presented by Yu et al. [140]. The QLSTM model enhances traditional LSTM networks by embedding VQCs, which optimize the weight parameters of the LSTM's gates, leading to better feature extraction and data compression. The study evaluates the QLSTM model's performance using data from five solar radiation observatories in China, comparing it with other models like SARIMA, CNN, RNN, GRU, and traditional LSTM. The results show that the QLSTM model consistently outperforms these models across various metrics, including RMSE, MAE, and R-Square, indicating its superior predictive accuracy. The QLSTM model's ability to handle different weather conditions and geographical locations further demonstrates its robustness and potential for broader applications in renewable energy forecasting.

Another work on solar irradiance forecasting discusses the development and comparison of hybrid quantum neural networks (HQNNs), with a focus on GPU-based platforms [60]. The authors highlight the design of a Hybrid Quantum Convolutional Neural Network (HQCNN) that integrates classical and quantum computing elements to enhance prediction accuracy. The study utilizes a dataset from Taiwan to optimize the network architecture through Bayesian Optimization, ensuring the model is well-suited for local conditions. The HQCNN is compared with traditional neural networks, demonstrating significant improvements in testing loss, which indicates better performance in predicting solar irradiance. The document also evaluates different platforms for developing HQNNs, including PennyLane, CUDA Quantum, and Torchquantum, with CUDA Quantum showing superior

performance in terms of speed and accuracy. The robustness of the HQCNN is tested under various simulated sensor fault scenarios, showcasing its ability to maintain performance despite noise and data loss. The study concludes that HQNNs, particularly when developed on CUDA Quantum, might offer substantial advantages over classical models, making them a promising approach.

Another quantum-enhanced technique in predicting solar irradiance has been recently proposed [130]. The authors focus on integrating quantum layers into traditional deep feedforward neural networks (FFNs) and compare these with a Bidirectional Long Short-Term Memory (BiLSTM) model, which serves as a benchmark. Data from NASA's POWER website, covering a decade, was used to train and validate the models. The data was scaled, and features were selected based on their relevance to solar irradiance prediction. After that, they developed quantum-enhanced FFNs, incorporating Parameterized Quantum Circuits (PQC) as quantum layers within the network. Their FFN models vary in complexity, with some having more quantum layers and/or different network architectures. The study also introduces a standalone PQC model, which uses quantum circuits exclusively for predictions. Hyperparameter tuning was conducted using the ADAM optimizer, and the models were trained over multiple epochs. The performance of each model was evaluated using metrics like Mean Absolute Error (MAE) and Root Mean Squared Error (RMSE). Results showed that hybrid models with quantum layers generally outperformed classical FFNs, with the FFN5L2Q hybrid model with two quantum layers standing out for its accuracy and efficiency. However, the standalone PQC model did not perform as well, indicating room for improvement in its design.

Very similar hybrid quantum machine learning models have been designed to improve the accuracy of direct photovoltaic (PV) power forecasting. A comprehensive study introduces three innovative models: Hybrid Quantum Neural Network (HQNN), Hybrid Quantum Long Short-Term Memory (HQLSTM), and Hybrid Quantum Sequence-to-Sequence (HQSeq2Seq) [113]. The HQNN model combines classical neural network layers with a quantum layer to predict solar power output for the next hour based on weather data from the previous 24 hours. Surprisingly, it outperforms its classical counterpart, the Multilayer Perceptron (MLP), by 41% in prediction accuracy while using fewer parameters. The HQLSTM model is a quantum-enhanced version of the classical LSTM, which is widely used for time series forecasting. The HQLSTM integrates quantum layers into each gate of the LSTM cell, resulting in a 40% improvement in prediction accuracy compared to the classical LSTM, despite having fewer parameters. The HQSeq2Seq model extends the capabilities of the classical Sequence-to-Sequence model by incorporating quantum layers. It allows for predicting PV power over arbitrary time intervals without needing prior weather data, making it highly versatile. The HQSeq2Seq model reduces prediction errors by 16% compared to its classical version. This study highlighted that hybrid quantum models not only provide more accurate forecasts but also perform better with limited datasets, making them suitable for scenarios where data collection is challenging.

Another study explores the potential of Quantum Long Short-Term Memory (QLSTM) models in improving solar power forecasting, comparing them with classical Long Short-Term Memory (LSTM) models [75]. Authors employed real-world solar data from plants in India and simulated data from the National Renewable Energy Laboratory (NREL) to ensure a comprehensive analysis. The research employs the PennyLane quantum machine learning framework, which integrates quantum and classical computing elements. The QLSTM model replaces traditional LSTM cells with variational quantum circuits. Both models are designed to handle the temporal dependencies in the data, but QLSTM uses quantum circuits to potentially offer a more nuanced understanding. The QLSTM model shows promising results, with faster convergence and better predictive accuracy compared to classical LSTM models. It achieves lower test losses and demonstrates a more stable learning process, indicating its potential. While QLSTM models show significant advantages, they also face challenges such as longer evaluation times due to the complexity of quantum simulations, as well as the need for further validation across diverse datasets and conditions.

Other than PV power generation, eolic power generation is the second source of RE by size and importance. As for solar irradiation, wind has a very aleatory behavior and this poses a great interest

in having solid forecasting to better control power generation and unit commitment. Hong, Ying-Yi et al. discuss a robust hybrid model that combines classical and quantum computing techniques to improve short-term wind speed forecasting [59]. This model integrates LSTM networks with QNNs to leverage the strengths of both classical and quantum computing. To ensure the model's robustness across different seasonal data variations, the Taguchi method is employed. This method systematically determines the optimal hyperparameters for the model, such as the number of LSTM layers and the type of quantum embedding, without relying on trial-and-error. The study finds that using IQP embedding with a quantum circuit depth of two yields the best results, making the model less sensitive to seasonal changes in wind speed data. The proposed hybrid model outperforms several existing models, including Random Forest, Support Vector Regression, and classical LSTM, particularly in handling the variability of wind speeds across different seasons.

A novel approach to forecasting power output from offshore wind farms using a hybrid quantum-classical machine learning model has been proposed [50]. The model utilizes a QNN as a feature extractor within an implicit variational approach. The QNN processes classical data by encoding it into quantum states, using amplitude encoding to map data into a higher-dimensional space. This involves a multi-layer PQC that applies various quantum gates to transform the input data. The extracted features are then fed into a Support Vector Regression (SVR) model, which is a classical machine learning technique used for regression tasks. The SVR is optimized using a grid search method to fine-tune its hyperparameters, ensuring the best possible prediction performance. The study evaluates the model's performance using data from a European offshore wind farm, spanning from 2010 to 2020. The model's effectiveness is measured using two metrics: Mean Absolute Error (MAE) and the coefficient of determination (R^2), indicating that the hybrid model successfully leverages quantum computing's capabilities to improve forecasting accuracy.

Another novel hybrid approach has been very recently presented by Hsu, Yu-Chao [61], as they introduce a model called Quantum Kernel-Based Long Short-Term Memory (QK-LSTM) for improving climate time-series forecasting. Even though they specifically focused on predicting the Air Quality Index (AQI), the model is designed to forecast other important weather data for indirect power generation like temperature, precipitation, irradiance or wind speed. The QK-LSTM model integrates quantum kernel methods into the LSTM framework by embedding classical data into high-dimensional quantum feature spaces. This allows the model to capture complex, nonlinear relationships in the data. The quantum kernel replaces traditional linear transformations in LSTM cells, enabling the model to handle intricate dependencies more effectively. Quantum kernel computations are performed using a specially designed quantum circuit that efficiently processes data in these high-dimensional spaces. The QK-LSTM model achieved higher predictive accuracy while using significantly fewer trainable parameters compared to classical LSTM models. This reduction in parameters decreases computational demands and reduces the risk of overfitting, leading to better generalization of the model. In addition, The QK-LSTM is designed to be scalable and practical for real-world applications. It supports hybrid quantum-classical implementations, making it feasible for deployment on current quantum hardware and in resource-constrained environments.

5.4. Grid Operations and Maintenance

5.4.1. Overview

In recent years, the energy and utilities sector has faced unprecedented challenges due to increasing demand, the integration of renewable energy sources and new actors, the so-called prosumers [38], and the need for efficient grid management. This variability can lead to fluctuations in power generation, making it difficult to maintain a stable and balanced grid. Additionally, renewable energy sources often lack the inertia provided by traditional synchronous generators, which can further aggravate stability issues. Faults in power systems can also have significant consequences, including equipment damage, power outages, and even cascading failures that can lead to widespread blackouts. Fault detection and fault diagnosis (FDD) involves identifying and diagnosing these faults to ensure the safe

and reliable operation of power systems. Hence early and accurate fault detection and diagnosis are crucial for several reasons:

- Increased safety and reliability: particularly important in high-risk environments such as power plants and transmission lines.
 - Reduced maintenance costs: enabling predictive maintenance, where potential faults can be identified before they cause significant damage. Reducing in turn maintenance costs and extending the lifespan of equipment. QML has the potential to further enhance predictive maintenance by improving the accuracy and speed of fault detection, leading to even greater cost savings and reduced downtime [26].
 - Efficiency: FDD can help improve the overall efficiency of power generation and distribution, by identifying and correcting inefficiencies in the system.
- Traditional FDD rely on various techniques, such as:
- Rule-based techniques: methods that use predefined rules and heuristics to identify faults based on specific patterns or thresholds in sensor data.
 - Model-based techniques: involve the systematic analysis of a system's anticipated behavior, which is grounded in physical and engineering principles. This anticipated behavior is then compared to the actual performance of the system. Deviations from this expectation could be identified as potential faults.

In the era of big data classical techniques have been evolved toward data-driven approaches employing machine and deep learning models, recent methods can be particularly effective in detecting complex or subtle faults that may not be captured by traditional methods. The volume and complexity of these data and problems could be overwhelming for traditional methods in the future. Although still in its early stages, Quantum Machine Learning (QML) is already showing promise in significantly improving these processes by leveraging its unique capabilities.

5.4.2. Key Studies

As said, the growing complexity of the energy grid makes it prone to faults and uncertain stability. In this picture fault analysis and diagnosis constitute an important aspect of electrical power systems and play a major role in handling severe failures caused by cascading effects of faults. Fault diagnosis tasks largely depend on feature extraction from the measured signals. The extracted feature characteristics directly contribute towards the effectiveness of fault detection and identification and classification accuracy of such fault diagnosis methods may be affected by the different features extracted from different feature extraction methods. In their work Akshay Ajagekar, Fengqi You (2021)[3] proposed a hybrid QC-based deep learning framework for fault diagnosis of electrical power systems that combine the feature extraction capabilities of conditional restricted Boltzmann machine with an efficient classification of deep networks. The idea is to overcome deep learning computational challenges by leveraging a QC-based training methodology. Specifically, the choice of CRBM has been made thanks to their ability of model temporal structures, making them a good choice for modeling multivariate time-series data. However, those complex model requires large training times and computation. The model architecture proposed by the authors is trained in a twofold approach:

- *Quantum generative training*: it initialize CRBM weights randomly and bias as zero vectors, then data and model expectations are computed by averaging the latent output variables and via quantum sampling respectively. Quantum sampling is performed on a quantum annealer, hence the problem should be formulated in such a way that is compatible with the QPU architecture. At every step of the training process the model parameters are updated via gradient ascent (mini batch fashion for stochasticity).
- *Discriminative training*: following generative training, discriminative training of the CRBM is performed. Data abstractions extracted from the CRBM are used to identify the state of the input measured data samples. The CRBM network with model parameters forms the first fully connected layer of the classification network. Those are already trained and will be fine-tuned

through this phase. Directed links between conditioning and hidden layers of the CRBM are treated as FFNN with a ReLU. On top of this, an additional fully connected layer is applied and finally, a sigmoid layer is used to predict class scores for each category.

The architecture has been tested by the authors on the IEEE 30 bus system, a widely used electrical power system for testing concepts by researchers which can be found on the IEEE power system test archive [23]. Performances of the architecture have been then compared to SOTA classical techniques such as ANN-based models [127] and DT [114], achieving comparable and, in some cases, even better results in terms both of Missed Detection Rate (MDR) and False Alarm Rate (FAR). Another interesting result regards the response time, the hybrid QC-CRBM framework showed a response time of 5 ms, same as DT model, while the ANN takes up to 10 ms for the classification of faults.

Studies on specific renewable source of energy has been made as of the case of the study by Glen Uehara et al.(2021)[135], in their work on “Quantum Neural Network Parameter Estimation for Photovoltaic Fault Detection”. The study investigates whether hybrid quantum-classical neural networks (QNNs) can provide enhanced accuracy and computational efficiency for PV fault detection in solar arrays, with the future idea of exploiting real-time data provided by smart monitoring devices (SMD). Authors have implemented a quantum computing based neural network in different versions to test which model have best performance in categorizing faults. This has been made in a hybrid fashion, where a quantum circuit is the hidden layer of a NN, and its parameters are optimized classically. The first QNN (named QNN1) has been designed with a *ZZFeatureMap*, while for the ansatz the variational form has been designed with *RealAmplitudes*, both from the Qiskit library. The other model called QNN2 has been designed with a *PauliFeatureMap* and the ansatz built-in *EfficientSU2* has been implemented, here the goal was to try to increase the accuracy with a fewer number of epochs, leveraging the ability of the feature map and ansatz in modeling more complex relationships. Both models have been tested on the NREL dataset[31] showing comparable results with respect to classical methods such as NN's.

The study from Camila Correa-Jullian et al. (2022)[26] explores quantum machine learning (QML) for wind turbine pitch fault detection using SCADA sensor data. Pitch system failures may account for up to 20% of total turbine downtime [79]. In this work, Quantum support vector machines (Q-SVMs) are compared with traditional machine learning (ML) models, including support vector machines (SVMs), random forests (RFs), and k-nearest neighbors (k-NN). While principal component analysis (PCA) and autoencoders (AE) are evaluated as feature reduction techniques. Results indicate that Q-SVMs perform competitively with classical methods and outperform RF and k-NN in certain feature settings. Although Q-SVMs require longer training times, advancements in quantum hardware may enhance their applicability in prognostics and health management (PHM).

Wind turbine systems (WTSs) rely on SCADA data for fault detection, that is a computer-based system that gathers and processes monitoring data from sensors, sampling every 1s but reporting a mean over 10min. There are two main problems to overcome in WTSs diagnostic with SCADA data:

- Most algorithms are unsupervised or semi-supervised, due the difficulty of acquiring robust and reliable labels from the systems [79].
- It is unclear which architecture and approach is better: SVM could be unstable in high dim, RF overfits easily, and DL models are highly complex and require huge amount of data to be trained on, hence can perform poorly under limited data availability regimes [39].

In this context Quantum computing offers new opportunities for improving classification tasks in PHM. To set up the framework of the case study, SCADA data from 2015 to 2019 were analyzed, including 251,164 temporal sensor entries and 337,448 alarms. The dataset underwent preprocessing, including handling missing values, standardizing sensor readings, and mapping alarm logs to sensor measurements. Feature reduction was performed using both PCA and AE. Those techniques have also shown denoising properties [129]. PCA retained 90% of the variance using 19 principal components, while AE utilized a nonlinear neural network to generate optimized feature representations. In principle using a linear activation should have same results of the PCA but a NN-Based AE with

nonlinear activation is expected to obtain a smaller and more accurate latent space of dimensions. Indeed, AE start to converge to a good MSE when around 15-16 features but it is also true that MSE does not provide a metric for understanding how much information is contained (as cumulative Variance for PCA). Moreover, AE does not allow to manage void entries, reducing the available data. Q-SVM models employed quantum kernels using both amplitude and angular encoding techniques as a feature map. For what concerns classical benchmarks, traditional SVMs with linear and radial basis function (RBF) kernels, RF, and k-NN were implemented, and model performance was assessed using a stratified k-fold cross-validation approach to ensure robustness and mitigate bias.

The study tested feature dimensionalities of 4, 8, 16, 19, and 32 to evaluate Q-SVM scalability. Amplitude and angular encoding were applied to lower dimensions, whereas for 16 and 32 features only amplitude encoding was used due to quantum circuit simulation constraints. For 19 features, zero-padding was employed to match hardware requirements.

Model training was conducted on classical hardware simulating quantum circuits. Each traditional ML model was trained 10 times independently, while Q-SVM models were trained 5 times due to computational constraints. Additionally, hyperparameter optimization was performed through grid search, and accuracy, F1-score, and standard deviation were used as evaluation metrics.

Q-SVMs demonstrated accuracy comparable to traditional ML models, particularly excelling with PCA-reduced features. AE-based reduction enhanced low-dimensional performance, while PCA was more effective at higher dimensions. Q-SVM performance varied with feature encoding, with angular encoding achieving the best results across all tested dimensions.

To formally assess the statistical difference between the models' performance and compare the ML models with the quantum kernels, a difference of means hypothesis test is performed for the best-performing data reduction configuration. The hypothesis test confirmed that Q-SVMs significantly outperformed RF and k-NN with statistical significance at $\alpha = 0.05$. However, Q-SVMs required substantially longer training time (~ 4 hours) compared to the one required for classical SVMs (seconds), due to the computational overhead of simulating quantum circuits on classical hardware.

A key analysis on power systems to ensure the resilience and reliability of modern interconnected grids is the so-called Transient Stability Assessment (TSA), which analyze the system ability to maintain synchronism when subjected to disturbances, such as short circuits or sudden loss of generation. Interconnected power systems are highly nonlinear, have different scale behaviors spatially and temporarily, and are also increasingly stochastic due to integration of renewable energy sources.

The work from Y. Zhou and P. Zhang (2021)[143] is the first attempt to unlock the potential of QML for power system TSA, with a focus in the design of an efficient and feasible qTSA circuit for near-term devices, but general enough to be exported in future noise free quantum computers. The key feature of qTSA, differently from classical machine learning techniques, is that the transient stability features in a Euclidean space are transformed to quantum states in a Hilbert space through a variational quantum circuit (VQC), which serves as a QNN to explicitly separate the stable and unstable samples. The framework employs a hybrid quantum-classical framework for QNN training. The parameterized circuit is executed on a quantum computer as the feedforward functionality of QNN, and parameter optimization is executed on a classical computer as the backpropagation functionality. Their architecture includes:

- Non-Gaussian feature encoding for flexible, nonlinear representation.
 - Gaussian quantum gates for efficient exploration of solution space.
 - Re-encoding layers to enhance expressivity.
 - A repetitive layered structure typical of VQC's.
- Authors also introduced a generalized quantum natural gradient to enable faster training.
- Authors have tested this framework on:
- Single-machine infinite-bus (SMIB) system, one of the most widely used test systems in power system research.
 - The two-area system, a benchmark system exhibiting both local and inter-area oscillation modes.

- Northeast Power Coordinating Council (NPCC) test system, a real Northeastern US power system[115].

The model has been tested on both a noise-free simulator and real quantum hardware to test resilience to noise on practical near-term devices.

Starting from simpler SMIB case study they demonstrated after training a classification accuracy exceeding 99% for both stable and unstable samples. A unique feature of qTSA is that it offers not only the stability classification results but also the fidelity of stability or instability. They have in fact demonstrated the faithful match between stable and unstable regions learned by the qTSA and samples based on analytical solutions, this allows the framework also to derive probabilities of system stability, an important insight for system operators and decision makers.

After that, test on the other larger scale systems has been carried on. For the small- and medium-scale power system, qTSA achieves high accuracy on both the training set (> 99%) and test set (> 98%). Even for the large-scale NPCC system, qTSA exhibits outstanding performance of 98% accuracy on the training set and 95% accuracy on the testing set.

Surprisingly, by running the framework on a real quantum device, the *ibmq_boeingen* QC, authors found that the qTSA performance was still of high quality, reporting a 1% decrease in accuracy even in the most challenging NPCC system. This experiment exhibits the effectiveness of qTSA on the near-term quantum devices and verifies the inherent resilience of qTSA against noisy quantum computing environments.

However, a larger scale qTSA circuit may fail to produce similar results, when testing on multiple configurations involving more qubits and accordingly deeper circuit depth. The accuracy on the real QC sharply deteriorates down to 52.02%, due to the noisy environment and quantum decoherence.

Various kinds of methods have been developed to cope with the problem of power system fault diagnosis. They can be generally divided into two main categories: data-driven and optimization-based methods. Until now, reported models belonged to the perimeter of the first category. In the latter, it is worth mentioning a work by Fei et al.(2024)[35] that proposes a quantum computing-based method for power system fault diagnosis using the Quantum Approximate Optimization Algorithm (QAOA). Their approach models the fault diagnosis problem as a Polynomial Unconstrained Binary Optimization (PUBO) and implements efficient quantum gate decomposition to reduce computational complexity. Key innovations include a symmetric equivalent decomposition (SED) method for multi-z-rotation gates and the use of small-probability event characteristics to lower qubit requirements. The method is validated on test systems, demonstrating competitive results with classical solvers while reducing computational time and resource usage. This study, although focused on combinatorial optimization in power systems, contributes to VQA research by advancing practical applications of QAOA in high-complexity domains.

Another work from Hangan et al. (2024)[51] explores the application of quantum machine learning for forecasting the stability of smart grids. They compare a classical Support Vector Machine (SVM) with a Variational Quantum Classifier (VQC) on a simulated dataset representing a 4-node smart grid architecture. While the classical SVM demonstrated superior performance in both training and testing phases, the study highlights challenges in quantum machine learning, such as sensitivity to noise and the complex fine-tuning of feature maps, ansatz, and optimizers. However, this work provides a foundational benchmark for integrating quantum methods into stability prediction in smart grids.

An interesting variational configuration called Quantum Variational Rewinding (QVR) is proposed by Jafari, Hossein and Aghababa, Hossein and Barati, Masoud [68]. The algorithm is designed for Time Anomaly Detection (TAD) using the Phasor Measurement Unit (PMU) dataset. This approach is specifically meant to leverage the capabilities of QPUs in the NISQ era, exhibiting a limited number of qubits and a shallow circuit depth. The QVR algorithm is introduced as a versatile tool that is not restricted to specific circuit structures, observables, or algorithmic elements, offering a high degree of adaptability. This flexibility allows for various algorithmic modifications to suit different appli-

cations. The primary focus of QVR is to analyze noisy single-variable electrical phasor data, which is crucial for real-time monitoring and anomaly detection in power distribution systems. The QVR algorithm is evaluated on IBM's Falcon r5.11H range of noise-affected superconducting transmon chips. The methodology emphasizes the need for a substantial number of measurements from simpler circuits, avoiding the complexity of intricate quantum circuits. The implementation of QVR involves integrating simultaneous QPU tasks with high-speed classical computing located close to the QPU. This integration is crucial to ensure that the total computation time remains within reasonable limits, making the algorithm practical for real-world applications.

Quantum computing techniques have been applied also for defect detection, helpful to prevent faults particularly, a study from Gbashi et al. (2024) [41] has implemented quantum capabilities into a classical framework, employing variational quantum circuits to enhance CNN model performance in diagnosing faults from high-dimensional vibration signals, which are typically costly to process using classical methods. Specifically, the study has been conducted on fault detection in wind turbine gearboxes, validating the proposed hybrid quantum-classical CNN (HQC-CNN) model using vibration signals sampled at 40 kHz from the intermediate-speed shaft bearing (IMSSB) of the gearbox of NREL's laboratory-scale wind turbine [83]. The proposed architecture is composed as follow:

- A classical section in which the classical data (28x28-sized image matrices) are processed into a CNN model with 5x5 kernel size and ReLu activation.
- A quantum section composed by a VQC with a ZZ-Feature-Map circuit for feature encoding layer and a Real-Amplitudes ansatz for the variational layer.
- A classical section in which the VQC parameters are optimized with classical optimization methods.

The study evaluates the model's performance across different optimizers, including adaptive moment estimation (Adam), stochastic gradient descent (SGD), and adaptive gradient algorithm (Adagrad) [138]. The Adam-based HQC-CNN model demonstrated superior efficiency, achieving optimal classification accuracy in 93 seconds, compared to 243 seconds for the SGD-based model. All models reach at least 99.2% accuracy in predicting gearbox health states albeit at different training epochs and computing times.

5.5. Finance for Sustainable Energy

5.5.1. Overview

The increasing global focus on climate change mitigation has led to the emergence of carbon markets as a key instrument for reducing greenhouse gas (GHG) emissions. These markets provide financial incentives for emissions reduction by enabling the trading of carbon credits, which represent the right to emit a certain amount of GHGs. Accurate forecasting of carbon prices and the estimation of carbon risk are crucial for the effective functioning of these markets and for informed decision-making by investors, policymakers, and other stakeholders. There many other financial operations related to sustainability and free energy market management that could benefit from ML and AI, such as profit-driven insights on energy trades, energy distribution scheduling, forecasting on national price trends, prediction of events' impact on the energy market, dynamic price incentivisation for REs usage, and so on.

5.5.2. Key Studies

A significant study conducted by Cao et al. (2023)[17] introduced the L-QLSTM, an enhanced version of a Long-Short Term Memory (LSTM) model, providing a specific Quantum Machine Learning framework for carbon price prediction. The basic idea is to leverage the advantage of LSTM architecture in temporal learning and utilize the expressive power that quantum entanglement offers. The authors based the paper on the work published by Chen et al. (2022)[22] in which they replaced the neural networks in different gates with Variational Quantum Circuits, finding that, under the constraint of similar number of parameters, this version learn faster and converges more stably than conventional

LSTM, especially when input data has complex temporal structures. In this framework, they brought some modification in the architecture:

- Linear layers before and after the VQC to extract feature representations. By compressing input features, linear layers reduce the number of qubits and considerably increase the learning ability of VQCs.
- The linear layer before the VQC's has shared parameters across all VQC's, to reduce parameters without losing a reduction in terms of parameters without losing too much accuracy in prediction.
- The variational layer from the original version employed CNOT operations to achieve entanglement. In this version the variational form of the VQCs is replaced by a strongly entangled controlled-Z quantum circuit. In principle, this should guarantee a stronger entanglement across qubits, giving them more expressivity.

For what concerns the parameters learning, VQC's parameters are optimized with a gradient based approach, similarly to classical neural networks. So given an output $f(x, \theta)$ with respect to the quantum parameters θ and input features x , the gradient of the VQC is computed with:

$$\nabla_{\theta} f(x, \theta) = \frac{1}{2} \left[f\left(x, \theta + \frac{\pi}{2}\right) - f\left(x, \theta - \frac{\pi}{2}\right) \right] \quad (18)$$

Then the loss could be minimized by the usual backpropagation of the gradient between VQC's and classical Neural Networks.

The empirical analysis of the framework has been applied to the biggest carbon trading market the European Climate Exchange (ECX), forecasting the settlement prices of the EUA continuous futures contracts that are traded in the ECX. Collected data are daily prices from Jan 2017 to Dec 2020, model predictive performance is evaluated based on Root Mean Square Error (RMSE) and Mean Absolute Error (MAE).

To verify whether L-QLSTM gives a performance improvement with respect to the basic QLSTM, and with respect to the classical LSTM, prediction performances are compared. Results shows a great improvement of the L-QLSTM against the base QLSTM and shows comparable performances with respect to the LSTM model.

Another interesting test was performed to see whether the SCZ variational form leads to better results compared to the variational form used in the original paper. Prediction errors on training and test datasets demonstrate a superior performance and improved stability of learning process of the former. Additionally, a parameter sensitivity analysis has been performed showing that model predictions are not sensitive to the selection of hidden dimension thus proving its robustness.

The other particularly promising area is risk estimation in carbon markets, where accurate forecasting of price fluctuations is essential for financial stability and environmental policy-making. Traditional risk estimation methods, such as time-series models, capital asset pricing models (CAPM) [102,111], and Monte Carlo simulations [54,57], are computationally intensive and very often rely on implicit assumptions about probability distributions that may only occasionally be reflected in reality [36]. In response to these challenges, quantum computing approaches have emerged as viable alternatives, offering novel methods for simulating risk distributions and enhancing measurement precision.

A recent study by Xiyuan Zhou et al. (2024) [142] explores this frontier by proposing a Quantum Conditional Generative Adversarial Network-Quantum Amplitude Estimation (QCGAN-QAE) framework for carbon market risk estimation. This approach leverages the advantages of quantum generative modeling and quantum measurement techniques, demonstrating notable improvements over classical risk estimation methods.

The European Union Emissions Trading System (EU ETS) serves as a case study in many carbon market analyses due to its significant role in global emissions regulation. Traditional risk metrics, such as Value-at-Risk (VaR) and Conditional Value-at-Risk (CVaR). They aim to quantify financial exposure, but their estimation is computationally expensive when using classical simulation-based approaches,

such as Monte Carlo methods. The inefficiency of these classical techniques stems from the need to sample large numbers of scenarios to achieve reliable risk estimates.

Quantum computing, particularly through quantum generative modeling, presents an alternative by leveraging the intrinsic parallelism of quantum states to encode and simulate probability distributions more efficiently, especially in the financial field [32]. The QCGAN-QAE framework introduces quantum generative adversarial networks (QGANs) to model future return distributions and employs quantum amplitude estimation (QAE) to measure risk metrics more effectively. The core innovation of the QCGAN-QAE framework lies in its twofold quantum approach: (i) employing Quantum Conditional Generative Adversarial Networks (QCGANs) for data generation and (ii) utilizing Quantum Amplitude Estimation (QAE) combined with binary search for efficient risk measurement.

QCGANs extend classical Conditional GANs (CGANs) by leveraging parameterized quantum circuits (PQCs) to generate simulated risk distributions. In this architecture, a quantum generator synthesizes return distributions based on historical data, while a classical discriminator evaluates their fidelity. One notable improvement in this study is the reordering of quantum entanglement and rotation layers, which enhances expressiveness and computational efficiency. Additionally, the introduction of a Quantum Fully Connected (QFC) layer allows for richer interactions among qubits, thereby improving the model's capacity to learn complex financial patterns. Once the QCGAN generates a simulated risk distribution, QAE is employed to measure VaR and CVaR efficiently. By utilizing binary search methods to speed up the process, QAE seems to improve upon classical Monte Carlo approaches, reducing the number of required simulations while maintaining precision.

The empirical analysis, based on EU ETS carbon price data from 2015 to 2020, demonstrates that QCGAN-QAE significantly improves risk estimation accuracy compared to classical methods. Results show that the QCGAN-QAE model generally outperforms historical simulation methods and classical CGANs, demonstrating improved accuracy in estimating VaR and CVaR in most cases. Although QCGAN-QAE performs well under normal market conditions, there are certain limitations in its ability to face extreme market risks.

Authors also analyzed the computational efficiency of the framework. Compared to classical CGANs, the QCGAN-QAE framework exhibits a 99.99% reduction in computational time, owing to the efficiency of quantum parallelism and the integration of binary search within QAE. Furthermore, the number of epochs required for model convergence is reduced by 90%, highlighting the potential of quantum generative models in risk assessment.

Further test has been performed on the QAE to demonstrate its advantages, proving better estimation accuracy and computational efficiency. Moreover, the QCGAN-QAE seems to be robust across different circuit depth, this could be an important feature for future scalability.

While the previous studies focus on Quantum Machine Learning approaches, alternative methodologies rely on pure optimization techniques, particularly in the formulation of Quantum Approximate Optimization Algorithm (QAOA) and Quadratic Unconstrained Binary Optimization (QUBO).

A study by Blenninger et al. (2024)[11] presents findings from the German Federal Ministry of Education and Research-funded project Q-GRID, which evaluates quantum optimization techniques in decentralized energy systems. The study examines how quantum optimization can improve demand-side management through personalized pricing strategies. The Discount Scheduling Problem (DSP) is formulated as a discrete optimization task where customers receive dynamic price incentives to shift energy usage based on CO₂ intensity forecasts. The optimization is encoded using binary variables and employs a QUBO-based approach. For benchmarking, the study compares classical optimization with quantum-hybrid techniques such as Gurobi [49], D-Wave's Leap Hybrid Solver [91], and Decomposition-based solvers [127] where problem instances are split into manageable subproblems solved independently. Specifically, Gurobi is the state-of-the-art classical solver, D-Wave Hybrid solver combines classical and quantum computing via annealing-based methods while Decomposition-based solvers have as core concept dividing the problem into manageable sub-problems that are then solved independently.

Results indicate that while Gurobi has good performances for small problem instances, hybrid quantum solvers show competitive performance as the problem scales, highlighting the potential of quantum techniques for large-scale energy scheduling.

Another application of machine learning for financial-oriented problems is reported by Andrés et al. [9]. Their work on quantum reinforcement learning using variational quantum algorithms counts three different use cases scenarios. The first, related to energy savings in buildings, has been already described in the previous section, along with the RL methodology and the model's architecture. The third use case assumes a scenario of scheduling control and real-time pricing in electrical vehicle (EV) charging stations. The main objective is to maximize the profit of the charging station, knowing that a proper scheduler must send customers to the service zone when they request service. With the same RL framework described for the first use case, the authors formulated specific-to-the-problem state space, action space and reward function.

The authors show that the quantum agent always converges to the same solution, which is the optimal policy.

On the other hand, the classical MLP agent is not able to learn the optimal policy. In this case, the quantum hybrid agent is also able to overcome its classical counterpart. This analysis is validated using a paired t-test with 95% confidence level, obtaining a p-value under 10^{-6} , which indicates that the performances of the quantum model significantly outperform its classical counterpart.

6. Analysis and Discussion

6.1. Technology Outlook

Quantum technology has advantages and challenges yet to be addressed. Research is ongoing to develop increasingly powerful and scalable hardware and software. While annealing is particularly suited to optimization problems, gate-based computation is more versatile and can be used for a wide range of applications. Emulated and simulated computation, on the other hand, offers a complementary approach that exploits quantum principles to foresee the performance of quantum computers once scaled at regime. From a physical point of view, the infrastructures of quantum installations are implemented using superconducting [77,105,107], photonics [1], topological [103,108], spin-based [16], neutral atom [55] or ion trap-based qubits[15], even if a real de facto technological standard has not yet established itself on the market and research is ranging in different directions. To answer the question of "if and when" these technologies are or will be applicable to real cases, it is necessary to introduce the concept of quantum advantage that describes the moment in which quantum computers surpass their classical equivalents in terms of performance on certain computational tasks.

Quantum advantage can be measured in terms of speed, scalability, energy efficiency or other relevant metrics. As a "rule of thumb", quantum advantage is mainly related to the number of "logical" qubits, that are the ones actually available for computation. Given the rise of QEC, as the number of qubits increases, in addition to a greater computational capacity, the guarantee of quantum computers' reliability also increases. This concept of quantum advantage is divided into three stages that reflect the degree of superiority. The first of the three stages is Quantum Utility, where quantum computers start to provide results in specific real and practical tasks or applications that cannot be matched with classical hardware. Even if the advantage of the quantum solution is not proven, it could still be considered relevant for practical purposes since it is something different from what you get with a classical computer. IBM among other hardware providers already believe they are in this stage [76]. Then there is Quantum Advantage, where quantum computers significantly surpass the capabilities of classical computers in a wider range of tasks or applications thanks to quantum error correction and reliable operations. When this level of advantage is achieved, quantum computers will demonstrate consistently better performance than classical computers on a variety of problems and scenarios. Quantum hardware authoritative players believe that this phase will be reached between 2029 and 2030 ([6,46,65,66,92,95,99,106,109]). Finally, there is Quantum Supremacy, that is, the highest level of advantage, in which quantum computers overwhelmingly and unequivocally surpass the capabilities

of classical computers via scalable and integrated hybrid hardware. The term "supremacy" was coined to describe the moment in which a quantum computer performs a certain operation in a time that is essentially impossible for any classical computer. For concrete applications, this will be reached once at least hundreds of thousands of qubits work in a stable and controlled environment. In the quantum roadmaps provided by hardware developers, the rise of this phase is currently scheduled for after 2030. At that point, once a stable logical qubit design has been achieved, efforts will focus on finding an efficient scaling method. Recent advancements suggest that this milestone may be approaching sooner than expected.

To conclude this brief excursus on quantum technologies, it is important to summarize the main reasons why it is important to start taking a concrete interest in Quantum Computing today. Competitive advantage and scarcity of computational resources are two sides of the same coin. On the one hand, the revolutionary potential that companies can seize by adopting quantum computing even early will be amplified by the natural scarcity of computational resources when installations will increasingly move towards "Advantage" and even more so in correspondence with "Supremacy". In this scenario, in fact, the demand for computational slots will increase much more rapidly than the corresponding supply by quantum computer constructs (see Figure 3). All this will cause the least ready companies, the "chronic followers", to wait: the time they will waste in line could also prove fatal for their business. Investing now in R&D or Pilot, and therefore in the concrete use of quantum in real cases, allows for further acceleration: the greater the interest and adoption by companies and research institutes, the greater the impetus to develop better hardware, software and algorithms. Finally, quantum computing is hoped to perform complex calculations in a more energy-efficient and therefore emission-efficient way, compared to classical computers. This will help reduce the environmental impact of intensive computational processing and promote more sustainable IT solutions. The conclusion is therefore to start to exploit quantum computing in a pragmatic way by intercepting the advantages it makes available today and preparing to seize the enormous opportunities that will be offered in a prospective way.




Figure 3. Impact of quantum computing technology on Predicted Temporal evolution of Offer-demand computational power curves. Once the Full Advantage is virtually reached (TP), Demand is expected to be greater than Offer. And even more so once Supremacy is achieved, Demand grows further. The portion of Demand that cannot be delivered due to Offer limits will remain on hold until the Offer curve reaches that corresponding level of capacity. The length of that time will be likely to dilate as technologies improve, causing companies that are not ready and waiting to lose important competitive advantage. As indicated by the red arrow, Supremacy is getting closer: several shreds of evidence suggest that technology is rapidly accelerating in various directions, and this moment is receding further and further behind with each detection.

6.2. Assessment Model for Innovation Management

In the following section, we will run an evaluation of the use cases previously introduced. This analysis, referred to as the Assessment Model for Innovation Management (AMIM), is designed to assess the suitability and benefits of each use case within the context of the industry’s current state and processes. It aids organizations in understanding the potential value and market readiness of a use case, providing a basis for strategic decision-making.

This epoch of the quantum era is characterized by strong uncertainty in terms of technology development and it is quite a challenge to foresee the applications that will benefit the most from quantum computing. Evaluations and estimations are very much dependent on strong assumptions mostly based on a mix of technology roadmaps’ predictions and theoretical algorithm advantages. Indeed, quantum computing workloads will be possibly integrated with classical architectures, at least at their first stage of development. This is particularly true for machine learning applications, where certain tasks will always be performed by classical hardware, and even more so in the case of variational quantum algorithms, which “per se” leverage quantum-classical integration. This inherent hybrid nature of quantum applications makes it difficult to really assess the technology’s performance but with heuristic-based real-world benchmarking. In parallel, limited research and development budgets prevent one from doing a blanket, indiscriminate search. Given this and that, understanding the best early opportunities of innovation with quantum computing is key for focusing human and time resources towards potential competitive advantages. That is what AMIM is meant for, no matter if it is for quantum machine learning or other cutting-edge innovation: the framework is posed to be unbiased with respect to the sector and the technology. This concept therefore perfectly suits the current context of our quantum machine learning use cases in the energy sector, offering opportunities of gathering resourceful insights for managing this new and risky field of innovation.

Several frameworks have inspired this concept, most of which focus on evaluating technologies rather than entire use cases. While developing our own framework, we found a few of these existing

frameworks particularly noteworthy. For instance, the Technology Readiness Level (TRL) and Technology Readiness Assessment (TRA) frameworks are widely used tools for assessing the maturity and readiness of technologies at various stages of development. TRL, initially developed by NASA for space programs, assesses the progression of a technology from basic principles to full-scale deployment, using a scale from 1 to 9. The TRA framework, on the other hand, evaluates the overall readiness of technologies by incorporating multiple criteria, such as technical, operational, and economic factors. In her work, Mialey Heder (2017) [64] critiques the increasing application of the TRL framework beyond its original context—specifically, its adoption in European Union policies and other areas where it was not initially intended. Heder argues that this growing use, without adequate rescaling to reflect the specific requirements of these new domains, risks undermining the sophistication and accuracy of the TRL framework's assessments. To address some of the limitations of the TRL framework, Kumari et al. (2022) [78] introduce the SMART (Software Maturity Assessment and Readiness Technique) framework developed as a two-dimensional model. While TRL focuses primarily on the technological readiness of a system, the SMART framework broadens the assessment scope by incorporating key factors such as quality evaluation, offering a more holistic approach. Thus, the SMART framework complements TRL by recognizing that a technology's readiness cannot be fully understood without considering both its maturity and quality.

Drawing from these foundations, we have developed our own framework, the Assessment Model for Innovation Management (AMIM) tailored to evaluate use cases on two dimensions:

- Readiness to Market
- Potential Benefit

6.2.1. Readiness to Market

This dimension measures how prepared a use case is to be deployed in the market, assessing factors like scaling, market compatibility, and implementation feasibility. The evaluation under this category focuses on whether the solution is viable and scalable in the industry today.

- Scalability: This KPI evaluates the potential for the use case to grow within the market. It considers whether it can scale as demand increases, handle larger user bases, and adapt to future needs. The criteria used to measure this KPI include the (i) *user growth potential*, which assesses its ability to accommodate increasing numbers of users, customers, or data, and (ii) *use case flexibility*, which evaluates the use case's capacity to integrate with new technologies and adjust to changing business environments.
- Market Compatibility: This KPI assesses how ready the present environment (e.g. society, stakeholders, technology, business, ecosystem) is to the use case. The criteria used to measure this KPI are (i) *customer readiness*, which evaluates the target audience's awareness and readiness to adopt the new use case and (ii) *technological infrastructure* which determines if the market has the required technology to support the use case
- Implementation Feasibility: This KPI assesses how ready the use case is to the present market. In other words it evaluates whether the use case integrates easily with existing systems and processes. The criteria used to measure this KPI include (i) *integration complexity*, which evaluates the number and complexity (customization requirements, compatibility, etc.) of integrations required with existing technologies, software, or hardware, and (ii) *compliance feasibility*, which assesses the ability to meet regulatory requirements, focusing on the ease and likelihood of achieving compliance.

6.2.2. Potential Benefit

This dimension measures the value the use case brings to the industry in terms of efficiency and competitive advantage. It evaluates the use case's ability to improve business outcomes and provide a long-term strategic advantage.

- **Impact on Efficiency:** This KPI measures the use case's potential to enhance operational efficiency. The criteria used to measure this KPI include (i) *cost reduction*, which evaluates the percentage reduction in operational or production costs post-implementation, (ii) *Return on Investment (ROI)*, which evaluates whether the benefits of the use case justify the investment required, determining if the use case is worthwhile in relation to the resources committed, and (iii) *productivity gains*, which measures the improvement in system productivity.
- **Criticality of the Problem:** This KPI measures the severity and importance of the problem being addressed. The more urgent or impactful the problem, the higher the benefit of solving it. The criteria used to measure this KPI include (i) *problem severity*, which assesses how serious and urgent the problem is for the target market, stakeholders, (ii) *market demand*, which evaluates the extent to which the market needs a solution to this problem, and (iii) *sustainability* which evaluates a use case's ability to promote long-term environmental health (resource consumption and waste), social well-being (community support), and ensure economic viability (financial stability).
- **Margin for Further Improvement:** This KPI measures how much the use case can be vertically and horizontally developed. The criteria used to measure this KPI include (i) *the development stage*, which determines the current stage of development of the use case, ranging from proposal stage to fully developed, thus reflecting the margin left for improvement. Additionally, (ii) *use case performance gaps* which helps identify any performance gaps in the present use case and therefore potential enhancements still necessary.

6.2.3. Results

The evaluation process involves scoring each use case across the three KPIs for each dimension described above, with scores that range from 1 to 4; 1 being weak, 2 low, 3 good and 4 high. The table (Table 4 in Appendix) organizes these scores, and the sum by dimension positions each use case on a graph (Figure 4). The x-axis represents Readiness to Market, while the y-axis represents Potential Benefit. The graph is divided into four quadrants:

- **Transformation Leaders (upper-right):** Use cases that are both market-ready and have high potential benefits.
- **Experimental Niche (lower-left):** Use cases that are not ready for market and offer low benefits.
- **Research Heavy Innovators (upper-left):** Use cases that are not market-ready but offer high potential benefits.
- **Emerging Niche (lower-right):** Use cases with low potential benefits but higher market readiness.

At the first sight, our analysis identifies the most essential use cases, which are called to be both market-ready and offer significant potential benefits. These include:

- ID 13: Dynamic Event Identification Using Phasor Measurement Units in Power Systems
- ID 9: Fault Diagnosis in Electrical Power Systems
- ID 5: Photovoltaic Power Systems

One would observe that the category "Grid Operations and Maintenance" stands for its more valuable use cases on average, which is coherent with the importance of such value chain elements in the market. Besides, there is still significant potential for improvement that this process can achieve with QML solutions. Conversely, financial use cases are prevalently experimental both in design and profit. Notably, although feasibility and technology maturity is pretty low for each use case of QML, the readiness to market dimension is quite decent for most. This is due to the high relevance of innovative features of scalability and response-to-market-demand for the energy industry. Hence, the industrial sector is one of the most interested actors in the adoption of quantum computing technologies.




Figure 4. Graphical Representation of AMIM. (For use case ID mapping see Table 1 in Appendix)

7. Conclusion

This study poses a systematic review of quantum machine learning early applications to the energy industry value chain. First, quantum computing and machine learning fundamentals have been provided for non-expert readers. Then, the study focuses on those quantum machine learning techniques that are hoped to solve real-world problems better than classical ones in the near-term future. These QML techniques have been found to be gate-based variational quantum algorithms, hybrid classical-quantum deep learning architectures, and Ising-like models. Reviewed key studies have been collected with specific inclusion and exclusion criteria based on the study’s focus. The resulting articles have been reported and analyzed. Results have proven promising future applications for QML in the energy industry. In particular, hybrid architectures could be the first step towards useful quantum computing in ML workloads, once integration feasibility and QPU performance issues have been solved. Given that, a novel general-purpose evaluation framework for managing innovation is proposed, called Assessment Model for Innovation Management (AMIM). This tool has been applied to QML in the energy industry value chain with interesting insights. This study shed light on possible future directions of QML use cases in such an important industrial sector. In the future, the presented framework could be used to tackle other industries, like financial services or automotive.

Acknowledgments: First and foremost, we are thankful to PwC Italy, particularly to our Data & AI leader Massimo Iengo for being the first to believe in our innovative journey into Quantum Computing. Not to forget all the people from the Data & AI team which have helped in making all of this work possible. We also wish to acknowledge the valuable financial support from Spoke 10 - ICSC - “National Research Centre in High Performance Computing, Big Data and Quantum Computing”, funded by European Union – NextGenerationEU. Notably, we express our sincere gratitude to Professor Paolo Cremonesi and Beatrice Goretti from Politecnico di

Milano for their essential support and encouragement since the very beginning of this entire Quantum Computing exploration initiative.

Appendix A

Table A1. Use Case Method’s Overview (Part 1)

ID	Method	Typology	SW Technology	HW Technology	Reported Benchmark
1	QSVM	Implicit	Not Specified	Not Specified	RNN, LSTM
	QNN	Data re-uploading	PennyLane, IBM Quantum Lab	IBM (various devices)	ARIMA, SARIMA, RNN, LSTM, GRU, Ensemble Learning
2	Hybrid RL	Hybrid	Qiskit	IBM Brisbane	MPC, DDPG, Lo-DDPG
3	Hybrid RL	Hybrid	Not Specified	Simulator	NN
4	QSVM	Explicit	Not Specified	Simulator	None
	QLSTM	Hybrid	PennyLane	Simulator	SARIMA, CNN, RNN, GRU, LSTM
	QNN	Data re-uploading	Qiskit	Simulator	SVR, XGBoost, GMDH
	Hybrid CNN	Hybrid	PennyLane, Torchquantum, CUDA	Simulator	CNN
	Hybrid NN	Hybrid	PennyLane	Simulator	RNN, LSTM
5	QNN, QLSTM, QSeq2Seq	Hybrid	PennyLane	Simulator	RNN, LSTM
	QLSTM	Hybrid	PennyLane	Simulator	LSTM
6	QLSTM	Hybrid	PennyLane	Simulator	RF, SVR, XGBoost, NAR, LSTM, LSTM AE
7	Hybrid QNN-SVR	Hybrid	PennyLane	Simulator	None
8	QK-LSTM	Implicit	Not Specified	Simulator	LSTM

Table A2. Use Case Method’s Overview (Part 2)

ID	Method	Typology	SW Technology	HW Technology	Reported Benchmark
9	Quantum Sampling for CRBM	Annealing	Ocean on Leap quantum cloud service	DWave 2000 QPU	NN, DT
10	QNN	Hybrid	Qiskit	Simulator	NN
11	QSVM	Implicit	Not Specified	Simulator	RF, k-NN, L-SVM, RBF-SVM
12	QNN	Data re-uploading	Qiskit	Simulator ibmq_boeligen QPU	None
13	QVR	Explicit	Qiskit	IBM Falcon r5.11H QPU	LSTM
14	VQC	Explicit	Not Specified	Simulator	SVM
15	Hybrid CNN	Explicit	Qiskit	Simulator	H-CNN versions
16	QLSTM	Hybrid	PennyLane, Pytorch	Simulator	QLSTM
17	Hybrid RL	Hybrid	Not Specified	Simulator	MLP
18	QCGAN + QAE	Data re-uploading	Qiskit	Simulator, IBM QPU	Historical simulation, CGAN, QCGAN

Table A3. Assessment Model for Innovation Management

ID	Scalability	Market Compatibility	Imple- mentation Feasibility	Tot. Readiness to Market	Impact on Efficiency	Criticality of the Problem	Margin for Further Improvement	Tot. Potential Benefit
1	4	2	1	7	3	3	3	9
2	4	1	2	7	4	3	1	8
3	4	3	2	9	1	1	1	3
4	4	4	2	10	1	3	2	6
5	3	4	2	9	3	3	2	8
6	4	3	1	8	2	2	3	7
7	2	3	3	8	2	2	2	6
8	4	3	3	10	1	1	2	4
9	4	4	4	12	3	4	3	10
10	3	4	1	8	2	1	2	5
11	2	4	2	8	1	1	2	4
12	4	2	1	7	4	4	4	12
13	4	4	1	9	3	4	4	11
14	3	2	1	6	3	3	3	9
15	2	3	3	8	1	1	3	5
16	3	2	2	7	1	2	2	5
17	1	1	2	4	1	1	1	3
18	3	3	1	7	2	2	3	7

References

1. H Aghaee Rad, T Ainsworth, RN Alexander, B Altieri, MF Askarani, R Baby, L Banchi, BQ Baragiola, JE Bourassa, RS Chadwick, et al. Scaling and networking a modular photonic quantum computer. *Nature*, pages 1–8, 2025.
2. A. Ahmad, D. Alahakoon, and I. Awan. A novel energy demand prediction strategy for residential buildings based on ensemble learning. *Energy*, 230:120–813, 2021.
3. Akshay Ajagekar and Fengqi You. Quantum computing based hybrid deep learning for fault diagnosis in electrical power systems. *Applied Energy*, 303:117628, 2021.
4. Akshay Ajagekar and Fengqi You. Variational quantum circuit based demand response in buildings leveraging a hybrid quantum-classical strategy. *Applied Energy*, 364:123244, 2024.
5. Tameem Albash and Daniel A. Lidar. Demonstration of a scaling advantage for a quantum annealer over simulated annealing. *Physical Review X*, 8(3):031016, 2018.
6. Alice and Bob. Quantum roadmap, 2025. Accessed: 2025-03-14.
7. A. Alonso, J. Perez, and M. Santos. Energy demand models for policy formulation: A comparative study of energy demand models. *Energy Policy*, 140:111–418, 2020.
8. Maria Andersson. Modeling electricity load curves with hidden markov models for demand-side management status estimation. *International Transactions on Electrical Energy Systems*, 27(3):e2265, 2017.
9. Eva Andrés, Manuel Pegalajar Cuéllar, and Gabriel Navarro. On the use of quantum reinforcement learning in energy-efficiency scenarios. *Energies*, 15(16):6034, 2022.
10. Jacob Biamonte, James Botterill, and et al. Quantum machine learning. *Nature*, 549:100–105, 2017.
11. Jonas Blenninger. Quantum optimization for the future energy grid: Summary and quantum utility prospects, 03 2024. To appear.
12. Denis Bokhan, Alena S Mastiukova, Aleksey S Boev, Dmitrii N Trubnikov, and Aleksey K Fedorov. Multi-class classification using quantum convolutional neural networks with hybrid quantum-classical learning. *Frontiers in Physics*, 10:1069985, 2022.
13. Léonard Bottou, Frank E. Curtis, and Jorge Nocedal. Optimization methods for large-scale machine learning. *SIAM Review*, 60(2):223–311, 2018.
14. George EP Box, Gwilym M Jenkins, Gregory C Reinsel, and Greta M Ljung. *Time series analysis: forecasting and control*. John Wiley & Sons, 2015.
15. Colin D Bruzewicz, John Chiaverini, Robert McConnell, and Jeremy M Sage. Trapped-ion quantum computing: Progress and challenges. *Applied physics reviews*, 6(2), 2019.
16. Guido Burkard, Thaddeus D Ladd, Andrew Pan, John M Nichol, and Jason R Petta. Semiconductor spin qubits. *Reviews of Modern Physics*, 95(2):025003, 2023.
17. Yuji Cao, Xiyuan Zhou, Xiang Fei, Huan Zhao, Wenxuan Liu, and Junhua Zhao. Linear-layer-enhanced quantum long short-term memory for carbon price forecasting. *Quantum Machine Intelligence*, 5:26, 2023.
18. Matthias C Caro, Elies Gil-Fuster, Johannes Jakob Meyer, Jens Eisert, and Ryan Sweke. Encoding-dependent generalization bounds for parametrized quantum circuits. *Quantum*, 5:582, 2021.
19. Marco Cerezo, Andrew Arrasmith, Ryan Babbush, Simon C Benjamin, Suguru Endo, Keisuke Fujii, Jarrod R McClean, Kosuke Mitarai, Xiao Yuan, Lukasz Cincio, et al. Variational quantum algorithms. *Nature Reviews Physics*, 3(9):625–644, 2021.
20. Marcos Cerezo and et al. Barren plateaus in quantum neural network training landscapes. *Nature Communications*, 12(1):1–12, 2021.
21. J. Cerrato et al. Quantum-inspired algorithms: A review and perspectives. *IEEE Transactions on Quantum Engineering*, 4:1–20, 2023.
22. Samuel Yen-Chi Chen, Shinjae Yoo, and Yao-Lung L. Fang. Quantum long short-term memory. In *ICASSP 2022 - 2022 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pages 8622–8626, 2022.
23. Richard D. Christie. Uw power systems test case archive. <http://labs.ece.uw.edu/pstca/>, 2018. Accessed: [Insert date of access here].
24. European Commission. Next generation eu. https://next-generation-eu.europa.eu/index_en. Accessed: 2023-10-05.
25. Don Coppersmith. Matrix inversion with quantum mechanics. In *Proceedings of the 35th Annual ACM Symposium on Theory of Computing*, pages 23–31. ACM, 1994.

26. Camila Correa-Jullian, Sergio Cofre-Martel, Gabriel San Martin, Enrique Lopez Droguett, Gustavo de Novaes Pires Leite, and Alexandre Costa. Exploring quantum machine learning and feature reduction techniques for wind turbine pitch fault detection. *Energies*, 15(8):Article 2792, 2022.
27. K. J. Daly and H. T. Hu. Fitting curves to data using nonlinear regression: A practical and nonmathematical review. *Creative Education*, 11(5):733–748, 2020.
28. Kumar Biswajit Debnath and Monjur Mourshed. Forecasting methods in energy planning models. *Renewable and Sustainable Energy Reviews*, 88:297–325, 2018.
29. C. Ding, F. Dong, and Z. Tang. Research progress on catalysts for the electrocatalytic oxidation of methanol. *ChemistrySelect*, 5(40):13318–13340, 2020.
30. Vivek Dixit, Raja Selvarajan, Tamer Aldwairi, Yaroslav Koshka, Mark A Novotny, Travis S Humble, Muhammad A Alam, and Sabre Kais. Training a quantum annealing based restricted boltzmann machine on cybersecurity data. *IEEE Transactions on Emerging Topics in Computational Intelligence*, 6(3):417–428, 2021.
31. A. Dobos. Pvwatts version 1 technical reference. Technical Report NREL/TP-6A20-60272, National Renewable Energy Laboratory, 2013. Online resource (8 pages).
32. Daniel J. Egger, Claudio Gambella, Jakub Marecek, Scott McFaddin, Martin Mevissen, Rudy Raymond, Andrea Simonetto, Stefan Woerner, and Elena Yndurain. Quantum computing for finance: State-of-the-art and future prospects. *IEEE Transactions on Quantum Engineering*, 1:1–24, 2020.
33. Y. Essam, A.N. Ahmed, R. Ramli, K.W. Chau, M.S.I. Ibrahim, M. Sherif, A. Sefelnasr, and A. El-Shafie. Investigating photovoltaic solar power output forecasting using machine learning algorithms. *Engineering Applications of Computational Fluid Mechanics*, 16:2002–2034, 2022.
34. Jeffrey Farhi, Edward Goldstone and Sam Gutmann. A quantum adiabatic evolution algorithm applied to random instances of an np-complete problem. *Science*, 292(5516):472–475, 2001.
35. Xiang Fei, Huan Zhao, Xiyuan Zhou, Junhua Zhao, Ting Shu, and Fushuan Wen. Power system fault diagnosis with quantum computing and efficient gate decomposition. *Scientific Reports*, 14:16991, 2024.
36. Zhen-Hua Feng, Yi-Ming Wei, and Kai Wang. Estimating risk for the carbon market via extreme value theory: An empirical analysis of the eu ets. *Applied Energy*, 99:97–108, 2012.
37. Maurizio Ferrari Dacrema, Fabio Moroni, Riccardo Nembrini, Nicola Ferro, Guglielmo Faggioli, and Paolo Cremonesi. Towards feature selection for ranking and classification exploiting quantum annealers. In *Proceedings of the 45th International ACM SIGIR Conference on Research and Development in Information Retrieval*, pages 2814–2824, 2022.
38. Walter Leal Filho, Laís Viera Trevisan, Amanda Lange Salvia, Janaina Mazutti, Thais Dibbern, Salvador Ruiz de Maya, Elvira Ferrer Bernal, João Henrique Paulino Pires Eustachio, Ayyoob Sharifi, María del-Carmen Alarcón-del Amo, and Iryna Kushnir. Prosumers and sustainable development: An international assessment in the field of renewable energy. *Sustainable Futures*, 7:100158, 2024.
39. Olga Fink, Qin Wang, Markus Svensén, Pierre Dersin, Wan-Jui Lee, and Melanie Ducoffe. Potential, challenges and future directions for deep learning in prognostics and health management applications. *Engineering Applications of Artificial Intelligence*, 92:103678, 2020.
40. Lucas Friedrich and Jonas Maziero. Avoiding barren plateaus with classical deep neural networks. *Physical Review A*, 106(4):042433, 2022.
41. Samuel M. Gbashie, Obafemi O. Olatunji, Paul A. Adedeji, and Nkosinathi Madushele. Hybrid quantum convolutional neural network for defect detection in a wind turbine gearbox. In *2024 IEEE PES/IAS PowerAfrica*, pages 01–06, 2024.
42. Gian Gentinetta, Arne Thomsen, David Sutter, and Stefan Woerner. The complexity of quantum support vector machines. *Quantum*, 8:1225, 2024.
43. Francisco Javier Gil Vidal and Dirk Oliver Theis. Input redundancy for parameterized quantum circuits. *Frontiers in Physics*, 8:297, 2020.
44. Vittorio Giovannetti, Seth Lloyd, and A. Nakadai. Quantum ram. *Physical Review Letters*, 100(16):160501, 2008.
45. Li-Hua Gong, Jun-Jie Pei, Tian-Feng Zhang, and Nan-Run Zhou. Quantum convolutional neural network based on variational quantum circuits. *Optics Communications*, 550:129993, 2024.
46. Google Quantum AI. Quantum ai roadmap, 2025. Accessed: 2025-03-14.
47. Lov K. Grover. A fast quantum mechanical algorithm for database search. *Proceedings of the 28th Annual ACM Symposium on Theory of Computing*, pages 212–219, 1996.
48. Emily Grumbling and Mark Horowitz. *Quantum Computing: Progress and Prospects*. National Academies Press, 2019.

49. Gurobi Optimization, LLC. *Gurobi Optimizer Reference Manual*, 2023.
50. Batuhan Hangun, Emine Akpinar, Murat Oduncuoglu, Oguz Altun, and Onder Eyecioglu. A hybrid quantum-classical machine learning approach to offshore wind farm power forecasting. In *2024 13th International Conference on Renewable Energy Research and Applications (ICRERA)*, pages 1105–1110. IEEE, 2024.
51. Batuhan Hangun, Onder Eyecioglu, and Oguz Altun. Quantum computing approach to smart grid stability forecasting. In *2024 12th International Conference on Smart Grid (icSmartGrid)*, pages 840–843, 2024.
52. Aram W. Harrow, Avinatan Hassidim, and Seth Lloyd. Quantum algorithms for fixed qubit architectures. *Physical Review Letters*, 103(15):150502, 2009.
53. Vojtěch Havlíček, Antonio D Córcoles, Kristan Temme, Aram W Harrow, Abhinav Kandala, Jerry M Chow, and Jay M Gambetta. Supervised learning with quantum-enhanced feature spaces. *Nature*, 567(7747):209–212, 2019.
54. Zhijian He. Sensitivity estimation of conditional value at risk using randomized quasi-monte carlo. *European Journal of Operational Research*, 298(1):229–242, 2022.
55. Loïc Henriët, Lucas Beguin, Adrien Signoles, Thierry Lahaye, Antoine Browaeys, Georges-Olivier Reymond, and Christophe Jurczak. Quantum computing with neutral atoms. *Quantum*, 4:327, 2020.
56. Zachary Holmes and et al. Connecting ansatz expressibility to gradient magnitudes and barren plateaus. *Quantum*, 5:430, 2021.
57. L. Jeff Hong, Zhaolin Hu, and Guangwu Liu. Monte carlo methods for value-at-risk and conditional value-at-risk: A review. *ACM Trans. Model. Comput. Simul.*, 24(4), November 2014.
58. T. Hong, S. Fan, and S. E. Kjaerbye. Probabilistic electric load forecasting: A tutorial review. *IEEE Transactions on Power Systems*, 35(3):2060–2070, 2020.
59. Ying-Yi Hong, Christine Joy E Arce, and Tsung-Wei Huang. A robust hybrid classical and quantum model for short-term wind speed forecasting. *IEEE Access*, 2023.
60. Ying-Yi Hong, Dylan Josh D Lopez, and Yun-Yuan Wang. Solar irradiance forecasting using a hybrid quantum neural network: A comparison on gpu-based workflow development platforms. *IEEE Access*, 2024.
61. Yu-Chao Hsu, Nan-Yow Chen, Tai-Yu Li, Kuan-Cheng Chen, et al. Quantum kernel-based long short-term memory for climate time-series forecasting. *arXiv preprint arXiv:2412.08851*, 2024.
62. Yi-Chung Hu. Electricity consumption prediction using a neural-network-based grey forecasting approach. *Journal of the Operational Research Society*, 68(10):1259–1264, 2017.
63. Tak Hur, Leeseok Kim, and Daniel K Park. Quantum convolutional neural network for classical data classification. *Quantum Machine Intelligence*, 4(1):3, 2022.
64. Mihály Héder. From nasa to eu: the evolution of the trl scale in public sector innovation. *The Innovation Journal*, 22(2):1–23, 2017.
65. IBM. Superconducting quantum roadmap, 2025. Accessed: 2025-03-14.
66. IonQ. Trapped ion quantum computing, 2025. Accessed: 2025-03-14.
67. Zuhaimy Ismail, Riswan Efendi, and Mustafa Mat Deris. Application of fuzzy time series approach in electric load forecasting. *New Mathematics and Natural Computation*, 11(03):229–248, 2015.
68. Hossein Jafari, Hossein Aghababa, and Masoud Barati. Quantum leaps: Dynamic event identification using phasor measurement units in power systems. In *2024 International Conference on Smart Grid Synchronized Measurements and Analytics (SGSMA)*, pages 1–6. IEEE, 2024.
69. M.A. Jallal, S. Chabaa, and A. Zeroual. A novel deep neural network based on randomly occurring distributed delayed pso algorithm for monitoring the energy produced by four dual-axis solar trackers. *Renewable Energy*, 149:1182–1196, 2020.
70. Sofiene Jerbi, Lukas J Fiderer, Hendrik Poulsen Nautrup, Jonas M Kübler, Hans J Briegel, and Vedran Dunjko. Quantum machine learning beyond kernel methods. *Nature Communications*, 14(1):1–8, 2023.
71. Sonika Johri. Bit-bit encoding, optimizer-free training and sub-net initialization: techniques for scalable quantum machine learning. *arXiv preprint arXiv:2501.02148*, 2025.
72. Abhinav Kandala and et al. Hardware-efficient variational quantum eigensolver for small molecules and quantum magnets. *Nature*, 549:242–246, 2017.
73. Muhammad Kashif, Muhammad Rashid, Saif Al-Kuwari, and Muhammad Shafique. Alleviating barren plateaus in parameterized quantum machine learning circuits: Investigating advanced parameter initialization strategies. In *2024 Design, Automation & Test in Europe Conference & Exhibition (DATE)*, pages 1–6. IEEE, 2024.
74. SS Kavitha and Narasimha Kaulgud. Quantum machine learning for support vector machine classification. *Evolutionary Intelligence*, 17(2):819–828, 2024.

75. Saad Zafar Khan, Nazeefa Muzammil, Salman Ghafoor, Haibat Khan, Syed Mohammad Hasan Zaidi, Abdulah Jeza Aljohani, and Imran Aziz. Quantum long short-term memory (qlstm) vs. classical lstm in time series forecasting: a comparative study in solar power forecasting. *Frontiers in Physics*, 12:1439180, 2024.
76. Youngseok Kim, Andrew Eddins, Sajant Anand, Ken Xuan Wei, Ewout Van Den Berg, Sami Rosenblatt, Hasan Nayfeh, Yantao Wu, Michael Zaletel, Kristan Temme, et al. Evidence for the utility of quantum computing before fault tolerance. *Nature*, 618(7965):500–505, 2023.
77. Morten Kjaergaard, Mollie E Schwartz, Jochen Braumüller, Philip Krantz, Joel I-J Wang, Simon Gustavsson, and William D Oliver. Superconducting qubits: Current state of play. *Annual Review of Condensed Matter Physics*, 11(1):369–395, 2020.
78. Archana Kumari, Stefan Schiffner, and Sandra Schmitz. Smart: a technology readiness methodology in the frame of the nis directive. *arXiv preprint arXiv:2201.00546*, 2022.
79. Kevin Leahy, Colm Gallagher, Peter O'Donovan, Ken Bruton, and Dominic T. J. O'Sullivan. A robust prescriptive framework and performance metric for diagnosing and predicting wind turbine faults based on scada and alarms data with case study. *Energies*, 11(7):Article 1738, 2018.
80. S. Li, C. Yan, and X. Liu. A review of data-driven building energy consumption prediction studies. *Applied Energy*, 236:797–809, 2019.
81. Wei Li, Peng-Cheng Chu, Guang-Zhe Liu, Yan-Bing Tian, Tian-Hui Qiu, and Shu-Mei Wang. An image classification algorithm based on hybrid quantum classical convolutional neural network. *Quantum Engineering*, 2022(1):5701479, 2022.
82. Zhaokai Li, Xiaomei Liu, Nanyang Xu, and Jiangfeng Du. Experimental realization of a quantum support vector machine. *Physical review letters*, 114(14):140504, 2015.
83. H Link, W LaCava, J van Dam, B McNiff, S Sheng, R Wallen, M McDade, S Lambert, S Butterfield, and F Oyague. Gearbox reliability collaborative project report: Findings from phase 1 and phase 2 testing. Technical report, National Renewable Energy Lab. (NREL), Golden, CO (United States), 06 2011.
84. Xia Liu, Geng Liu, Hao-Kai Zhang, Jiaxin Huang, and Xin Wang. Mitigating barren plateaus of variational quantum eigensolvers. *IEEE Transactions on Quantum Engineering*, 2024.
85. Seth Lloyd, Maria Schuld, Aroosa Ijaz, Josh Izaac, and Nathan Killoran. Quantum embeddings for machine learning. *arXiv preprint arXiv:2001.03622*, 2020.
86. Yunpu Ma, Volker Tresp, Liming Zhao, and Yuyi Wang. Variational quantum circuit model for knowledge graph embedding. *Advanced Quantum Technologies*, 2(7-8):1800078, 2019.
87. R. Malshe, S. R. Deshmukh, and A. Khandare. Electrical load forecasting models: A critical systematic review. *Renewable and Sustainable Energy Reviews*, 131:109–974, 2020.
88. M.J. Mayer. Benefits of physical and machine learning hybridization for photovoltaic power forecasting. *Renewable and Sustainable Energy Reviews*, 168:112772, 2022.
89. Jarrod R. McClean and et al. The theory of variational hybrid quantum-classical algorithms. *New Journal of Physics*, 18(2):023023, 2016.
90. Jarrod R. McClean, Raphael Sweke, Ignacio Rungger, and et al. Cost function dependent barren plateaus in shallow parametrized quantum circuits. *Physical Review A*, 98(4):042324, 2018.
91. Catherine McGeoch, Pau Farre, and William Bernoudy. D-wave hybrid solver service + advantage: Technology update. Technical report, D-Wave Systems Inc., 2020.
92. MeetIQM. Technology roadmap, 2025. Accessed: 2025-03-14.
93. Ali Danandeh Mehr, Farzaneh Bagheri, and Mir Jafar Sadegh Safari. Electrical energy demand prediction: A comparison between genetic programming and decision tree. *Gazi University Journal of Science*, 33(1):62–72, 2020.
94. Riccardo Mengoni and Alessandra Di Pierro. Kernel methods in quantum machine learning. *Quantum Machine Intelligence*, 1(3):65–71, 2019.
95. Microsoft. Topological quantum roadmap, 2025. Accessed: 2025-03-14.
96. Lorenzo Moro and Enrico Prati. Anomaly detection speed-up by quantum restricted boltzmann machines. *Communications Physics*, 6(1):269, 2023.
97. Mounica Nutakki, Suprabhath Koduru, Srihari Mandava, et al. Quantum support vector machine for forecasting house energy consumption: a comparative study with deep learning models. *Journal of Cloud Computing*, 13(1):1–12, 2024.
98. Victor Oliveira Santos, Felipe Pinto Marinho, Paulo Alexandre Costa Rocha, Jesse Van Griensven Thé, and Bahram Gharabaghi. Application of quantum neural network for solar irradiance forecasting: A case study using the folsom dataset, california. *Energies*, 17(14):3580, 2024.

99. Pasqal. Our technology and quantum roadmap, 2025. Accessed: 2025-03-14.
100. Taylor L Patti, Khadijeh Najafi, Xun Gao, and Susanne F Yelin. Entanglement devised barren plateau mitigation. *Physical Review Research*, 3(3):033090, 2021.
101. Adrián Pérez-Salinas, Alba Cervera-Lierta, Elies Gil-Fuster, and José I Latorre. Data re-uploading for a universal quantum classifier. *Quantum*, 4:226, 2020.
102. Le Tan Phuoc, Kee S. Kim, and Yingcai Su. Reexamination of estimating beta coefficient as a risk measure in capm. *The Journal of Asian Finance, Economics and Business*, 5(1):11–16, February 28 2018.
103. Elsa Prada, Pablo San-Jose, Michiel WA de Moor, Attila Geresdi, Eduardo JH Lee, Jelena Klinovaja, Daniel Loss, Jesper Nygård, Ramón Aguado, and Leo P Kouwenhoven. From andreev to majorana bound states in hybrid superconductor–semiconductor nanowires. *Nature Reviews Physics*, 2(10):575–594, 2020.
104. John Preskill. Quantum computing in the nisq era and beyond. *Quantum*, 2:79, 2018.
105. Harald Putterman, Kyungjoo Noh, Connor T Hann, Gregory S MacCabe, Shahriar Aghaeimeibodi, Rishi N Patel, Menyong Lee, William M Jones, Hesam Moradinejad, Roberto Rodriguez, et al. Hardware-efficient quantum error correction via concatenated bosonic qubits. *Nature*, 638(8052):927–934, 2025.
106. Quantinuum. Accelerated roadmap to universal fault-tolerant quantum computing by 2030, 2025. Accessed: 2025-03-14.
107. Google AI Quantum. Suppressing quantum errors by scaling a surface code logical qubit. *Nature*, 614(7949):676–681, 2023.
108. Microsoft Azure Quantum, Morteza Aghaee, Alejandro Alcaraz Ramirez, Zulfi Alam, Rizwan Ali, Mariusz Andrzejczuk, Andrey Antipov, Mikhail Astafev, Amin Barzegar, Bela Bauer, et al. Interferometric single-shot parity measurement in inas–al hybrid devices. *Nature*, 638(8051):651–655, 2025.
109. Quera. Quantum roadmap, 2025. Accessed: 2025-03-14.
110. Manuel S Rudolph, Jacob Miller, Danial Motlagh, Jing Chen, Atithi Acharya, and Alejandro Perdomo-Ortiz. Synergy between quantum circuits and tensor networks: Short-cutting the race to practical quantum advantage. *arXiv preprint arXiv:2208.13673*, 2022.
111. Anna Rutkowska-Ziarko, Lesław Markowski, Christopher Pyke, and Saqib Amin. Conventional and downside capm: The case of london stock exchange. *Global Finance Journal*, 54:100759, 2022.
112. Ashkan Safari and Mohammad Ali Badamchizadeh. Neuroquman: Quantum neural network-based consumer reaction time demand response predictive management. *Neural Computing and Applications*, 36(30):19121–19138, 2024.
113. Asel Saginalieva, Stefan Komorniyk, Arsenii Senokosov, Ayush Joshi, Alexander Sedykh, Christopher Mansell, Olga Tsurkan, Karan Pinto, Markus Pflitsch, and Alexey Melnikov. Photovoltaic power forecasting using quantum machine learning. *arXiv preprint arXiv:2312.16379*, 2023.
114. S.R. Samantaray. Decision tree-based fault zone identification and fault classification in flexible ac transmissions-based transmission line. *IET Generation, Transmission & Distribution*, 3:425–436, 2009.
115. Peter W. Sauer, M. A. Pai, and Joe H. Chow. *Power System Dynamics and Stability: With Synchrophasor Measurement and Power System Toolbox*. Wiley-IEEE Press, 2017.
116. Bernhard Schölkopf, Ralf Herbrich, and Alex J Smola. A generalized representer theorem. In *International conference on computational learning theory*, pages 416–426. Springer, 2001.
117. Maria Schuld and Nathan Killoran. Quantum machine learning in feature hilbert spaces. *Physical review letters*, 122(4):040504, 2019.
118. Maria Schuld, Ryan Sweke, and Johannes Jakob Meyer. Effect of data encoding on the expressive power of variational quantum-machine-learning models. *Physical Review A*, 103(3):032430, 2021.
119. Alessandro Sebastianelli, Daniela Alessandra Zaidenberg, Dario Spiller, Bertrand Le Saux, and Silvia Liberata Ullo. On circuit-based hybrid quantum neural networks for remote sensing imagery classification. *IEEE Journal of Selected Topics in Applied Earth Observations and Remote Sensing*, 15:565–580, 2021.
120. Makhmisa Senekane, Benedict Molibeli Taele, et al. Prediction of solar irradiation using quantum support vector machine learning algorithm. *Smart Grid and Renewable Energy*, 7(12):293, 2016.
121. M. Shahid, S. Yahya, and M. A. Khan. Computational intelligence approaches for energy load forecasting in smart energy management grids: State of the art, future challenges, and research directions. *Energy Reports*, 7:1022–1035, 2021.
122. Shai Shalev-Shwartz and Shai Ben-David. *Understanding machine learning: From theory to algorithms*. Cambridge university press, 2014.
123. 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.

124. Jimann Shin, Jiakun Chen, and Lilianna Solnica-Krezel. Efficient homologous recombination-mediated genome engineering in zebrafish using tale nucleases. *Development*, 141(18):3807–3818, 2014.
125. Peter W. Shor. Scheme for reducing decoherence in a quantum computer. *Physical Review A*, 52(4):R2493–R2496, 1995.
126. Peter W. Shor. Algorithms for quantum computation: discrete logarithms and factoring. *Proceedings of the 35th Annual ACM Symposium on Theory of Computing*, pages 124–134, 1997.
127. Kleber Silva, B.A. Souza, and N.S.D. Brito. Fault detection and classification in transmission lines based on wavelet transform and ann. *Power Delivery, IEEE Transactions on*, 21:2058 – 2063, 11 2006.
128. Andrea Skolik, Jarrod R McClean, Masoud Mohseni, Patrick Van Der Smagt, and Martin Leib. Layerwise learning for quantum neural networks. *Quantum Machine Intelligence*, 3:1–11, 2021.
129. Adrian Stetco, Fateme Dinmohammadi, Xingyu Zhao, Valentin Robu, David Flynn, Mike Barnes, John Keane, and Goran Nenadic. Machine learning methods for wind turbine condition monitoring: A review. *Renewable Energy*, 133:620–635, 2019.
130. Mushrafi Munim Sushmit and Islam Mohammed Mahbubul. 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.
131. Yudai Suzuki, Hiroshi Yano, Qi Gao, Shumpei Uno, Tomoki Tanaka, Manato Akiyama, and Naoki Yamamoto. Analysis and synthesis of feature map for kernel-based quantum classifier. *Quantum Machine Intelligence*, 2:1–9, 2020.
132. Sohum Thakkar, Skander Kazdaghli, Natansh Mathur, Iordanis Kerenidis, André J Ferreira-Martins, and Samurai Brito. Improved financial forecasting via quantum machine learning. *Quantum Machine Intelligence*, 6(1):27, 2024.
133. Y. Tian, Q. Wu, and Z. Zhang. Multiscale stochastic prediction of electricity demand in smart grids using bayesian networks. *IEEE Transactions on Smart Grid*, 13(1):186–197, 2022.
134. Alexandros Tzanetos and Georgios Dounias. An application-based taxonomy of nature inspired intelligent algorithms. *Management and Decision Engineering Laboratory (MDE-Lab) University of the Aegean, School of Engineering, Department of Financial and Management Engineering, Chios*, 2019.
135. Glen Uehara, Sunil Rao, Mathew Dobson, Cihan Tepedelenlioglu, and Andreas Spanias. Quantum neural network parameter estimation for photovoltaic fault detection. In *2021 12th International Conference on Information, Intelligence, Systems & Applications (IISA)*, pages 1–7, 2021.
136. Walter Vinci and Daniel A. Lidar. Non-stoquastic hamiltonians in quantum annealing via geometric phases. *npj Quantum Information*, 3(1):1–7, 2017.
137. Nan Wei, Changjun Li, Xiaolong Peng, Fanhua Zeng, and Xinqian Lu. Conventional models and artificial intelligence-based models for energy consumption forecasting: A review. *Journal of Petroleum Science and Engineering*, 181:106–187, 2019.
138. Jiabin Yang and Qiang Long. A modification of adaptive moment estimation (adam) for machine learning. *Journal of Industrial and Management Optimization*, 20(7):2516–2540, 2024.
139. Yordan S Yordanov, Vasileios Armaos, Crispin HW Barnes, and David RM Arvidsson-Shukur. Qubit-excitation-based adaptive variational quantum eigensolver. *Communications Physics*, 4(1):228, 2021.
140. Yunjun Yu, Guoping Hu, Caicheng Liu, Junjie Xiong, and Ziyang Wu. Prediction of solar irradiance one hour ahead based on quantum long short-term memory network. *IEEE Transactions on Quantum Engineering*, 4:1–15, 2023.
141. Shuxiao Zhang, Teng Yang, and Shijun Zhan. Quantum ram with $o(1)$ query complexity. *Physical Review Letters*, 121(3):030501, 2018.
142. Xiyuan Zhou, Huan Zhao, Yuji Cao, Xiang Fei, Gaoqi Liang, and Junhua Zhao. Carbon market risk estimation using quantum conditional generative adversarial network and amplitude estimation. *Energy Conversion and Economics*, 5(4):193–210, 2024.
143. Yifan Zhou and Peng Zhang. Noise-resilient quantum machine learning for stability assessment of power systems. *IEEE Transactions on Power Systems*, 38(1):475–487, 2023.

Disclaimer/Publisher’s Note: The statements, opinions and data contained in all publications are solely those of the individual author(s) and contributor(s) and not of MDPI and/or the editor(s). MDPI and/or the editor(s) disclaim responsibility for any injury to people or property resulting from any ideas, methods, instructions or products referred to in the content.