

Uncertainty-Aware Quantum Machine Learning for Making Reliable Decision-Making in Noisy Environments

Manimaran R¹, Sathiyajithrey Loganathan², Justin Suresh³, Maruthu Kannan⁴

^{1,2,3,4} Kalasalingam Academy of Research and Education, India

Received Date: 6 April 2026

Revised Date: 17 April 2026

Accepted Date: 28 April 2026

Abstract

Quantum Machine Learning (QML) represents a new paradigm that combines quantum computing with classical machine learning to address complex computational problems. Yet, the effective utilization of QML systems is limited because of the noise naturally arising from quantum hardware and the uncertainty in data and model predictions. Uncertainty-aware QML: improving reliability and robustness in noise-prone environments (Reprint from Plops ONE)† — Introduction In this work we investigate uncertainty-aware quantum-machine learning (QML), with a specific emphasis on how to improv. Quantum systems are prone to DE coherence, gate errors and interaction with the environment, inducing noise that causes uncertainty in computational results. Finally, given that classical data used to train QML models is typically noisy by nature, the errors are compounded. This highlights the demand of uncertainty-aware frameworks that can capture and mitigate both quanta as well as classical uncertainties. In this paper we gather an extensive review of uncertainty embedding techniques: Probabilistic methods, Bayesian inference and hybrid quantum-classical methods. The overview also covers noise-aware architectures like quantum neural networks (QNNs), and quantum support vector machines (QSVM) or Variational quantum circuits. Additionally, the study examines sophisticated approaches including error mitigation, uncertainty quantification as well as robust optimization methods to improve decision-making reliability. The practical significance of uncertainty-aware QML is illustrated by examining applications in healthcare diagnostics, cybersecurity, finance, and autonomous systems. The paper wraps up with important research bottlenecks and future directions, including scalable quantum hardware, better noise modelling and explainable AI integration. This work contributes to the development of trustworthy QML systems that can make reliable decisions in practice by addressing uncertainty and noise jointly.

Keywords

Quantum Machine Learning, Uncertainty Quantification, Noisy Quantum Systems, NISQ, Quantum Neural Networks, Robust AI, Decision Making Hybrid Models

Introduction

Quantum computing, on the other hand, alters the very fabric of this computational paradigm; it uses quantum properties such as superposition and entanglement to execute tasks whose complexity is intractable for classical computers. While classical bits are in either a state of 0 or 1, quantum bits (quits) can be in multiple states at once thereby allowing them to do computations at the same time. Such methods permit quantum systems to approach certain categories of problems—like optimization, cryptography and simulation—in a more computationally efficient manner than classical methods. This paradigm is called Quantum Machine Learning (QML) which is a rapidly growing field that uses the benefits of quantum computation to improve learning algorithms.

QML has shown great potential in contexts like feature space transformations, kernel methods and Variational optimization. For accelerating pattern recognition and classification tasks, several algorithms have been proposed, including Quantum Support Vector Machines (QSVMs), Variational Quantum Circuits (VQCs) as well as Quantum Neural Networks (QNNs). Nonetheless, in spite of such theoretical progress, the implementation of QML is still limited by various issues — most fundamentally noise and uncertainty that are intrinsic to quantum systems.

Currently, quantum computing is in the era of Noisy Intermediate-Scale Quantum (NISQ), which features noisy quantum devices with a limited number of quits. Quits are notoriously delicate and quickly lose their quantum state through DE coherence due to interactions with the environment. Quantum gate operations are not error-free, and measurements introduce additional probabilistic errors over and above the quantum gates. All these



imperfections together produce quantum noise, which strongly degrades the fidelity and precision of quantum computations. Thus, QML models generated bins devices tend to have a high degree of variability in their outputs.

At the same time uncertainty from several sources also challenges classical machine learning systems. Some training data has noise; sometimes, results are based on measurements that can be imprecise for various reasons (mismeasurements, missing values), while in real-life processes there is always some variation for biological and physical reasons. This sort of uncertainty is called aleatoric – which refers to randomness in the data itself, and cannot be reduced by simply obtaining additional data. Epistemic uncertainty, on the other hand, stems from model limitations (e.g. This can occur due to a shortage of trained data or failing assumptions regarding the probabilistic distribution of underlying data). While aleatoric uncertainty cannot be eliminated, epistemic uncertainty can be decreased with better modelling and more data.

Quantum noisiness and classical uncertainty jointly induce a double trouble for QML systems. Quantum noise randomness is based on flip at the hardware level, whereas classical uncertainties interfere with the learning and model predications. The interplay of those two forms of uncertainty is clearly disruptive for the QML models and results in erroneous determinations. This problem can be especially dangerous in high-stake applications including but not limited to health care diagnostic tools, financial forecasting, cyber security, and autonomous systems where wrong predictions could have devastating consequences.

In response to these challenges, the field of uncertainty-aware QML has arisen as a potential key research direction. Uncertainty-aware QML explicitly models, quantifies and manages uncertainty across the entire machine learning pipeline. These models also incorporate probabilistic methods and statistical inference techniques which can be used not just to make predictions, but also to determine confidence, associated with the prediction. It is very useful as it helps the users to determine how reliable the outputs of a model are before making a decision-making.

An important strategy in uncertainty-aware QML is the combination of methods for Bayesian inference with quantum algorithms. Bayesian frameworks make it possible to represent uncertainty in terms of model parameters and predictions, which in turn makes adaptive learning among uncertain environments possible. In parallel, hybrid quantum-classical models have attracted interest in the ability to leverage the advantages of quantum computing and robust classical optimization techniques. Such hybrid methods may obviate the shortcomings of contemporary quantum hardware whilst exploiting quantum advantages for appropriate computational tasks.

Some recent work suggests that approaching quantum hardware noise and classical data uncertainty together as interconnected problems can be more advantageous instead of considering them separately. Approaches to making QML models more robust, including error mitigation [18], noise-aware training [19], and robust optimization [20] have been proposed. The success of such methods in accomplishing reliable predictions in noisy environments is also aided by advances on uncertainty quantification methodologies (e.g. Monte Carlo sampling, Variational inference and ensemble learning).

In this paper, we aim to provide an extensive overview of uncertainty-aware Quantum Machine Learning that can assist stakeholders with reliable decision-making in noisy settings. It investigates sources of noise and uncertainty, presents advanced techniques for quantifying uncertainty and mitigating noise, and tests QML models performance on real-world problems. The paper also identifies practical use cases and suggests avenues for future research, such as the creation of scalable quantum hardware and incorporation into explainable AI paradigms.

With its integrated handling of quantum and classical uncertainties, uncertainty-aware QML is a fundamental stepping stone toward confidence in quantum AIs that can function well outside the lab.




FIGURE 1: Types of Uncertainty in Machine Learning: Aleatoric And Epistemic Components

Literature Review and Underpinning of Concepts

Quantum Machine Learning (QML) has become a fast growing area of research bridging the gap between Quantum Computing and Machine Learning, with the goal of improving computational efficiency and learning

through quantum mechanical principles. This initial work in QML was mainly concerned with speeding up classical algorithms using quantum mechanisms, such as for linear algebra: e.g., the inversion of matrices and the estimation of eigenvalues. These basic results revealed the potential for quantum systems to provide exponential or polynomial speedups for some problems, generating compelling interest in new classes of learning models that are more powerful because they are enhanced with quantum mechanics.

One of the most salient advances in QML has been based around Variational Quantum Algorithms (VQAs)—parameterized quantum circuits, the parameters of which are tuned using classical feedback loops. These algorithms are well-suited for existing quantum hardware, referred to as the NISQ regime (Noisy Intermediate-Scale Quantum). They operate under a regime in which quantum devices are small and error-prone, so that today it is not feasible to have fully fault-tolerant quantum computation. To address this limitation, a new family of techniques called VQAs perform hybrid quantum-classical optimization, where the output from quantum circuits is used and classical optimizers update parameters in an iterative manner based on the minimized cost function. This has thus allowed for QML models such as classification and generative tasks to be conducted with responsible implementation even in the face of hardware limitations.

Nevertheless, quantum noise due to DE coherence, gate and measurement errors can largely affect the efficacy of QML models. One such process, DE coherence specifically defines the loss of quantum information over time that occurs when groups of qubits interact with external noise sources. The presence of many noise factors not only lowers the computational accuracy but also creates a lot of variability in predictions made by the model, which is a key obstacle to real-world deployments. Research in this field has established a general result that noise can corrupt quantum states and deteriorate model performance; notably for deep quantum circuits when the number of parameters become large.

In order to resolve these issues, noise-aware QML methods are devised that take the hardware imperfections into account during the modelling and training phases. Quantum Neural Networks (QNNs) are models that have been modified to use shallow circuit architectures, which allow them to be less susceptible to noise while retaining their expressive power [4]. Hybrid quantum-classical models, which leverage both paradigms, have also become popular. In such systems, quantum circuits are used for feature encoding and transformation whilst classical neural networks perform the optimization and decision aspect of learning. This synergy enhances robustness and enables models to perform satisfactorily despite the presence of noise. Meanwhile, it was established a fruitful area of focus on uncertainty aware machine learning in classical domains parallel with the improvements in QML. Common machine learning models are deterministically mapped functions providing no confidence margins on their output, hard to use in critical applications. In order to overcome this limitation, many probabilistic methods such as Bayesian inference and Gaussian processes have been proposed. Such methods allow models to measure uncertainty by modelling predictions as distributions instead of hard values. This ability is important because it can differentiate between aleatoric and epistemic uncertainty, which respectively stem from randomness of the data or a lack of knowledge about the model.

This brings the approach in noise-aware QML closer to those investigated for uncertainty-aware machine learning giving birth to uncertainty-aware QML frameworks. Both of these frameworks try to bridge quantum noise and classical uncertainty into one model. Researchers can improve the interpretability and reliability of quantum ML (QML) systems using probabilistic reasoning in our quantum algorithms. An example is Bayesian quantum models which provide an explicit treatment of uncertainty in circuit parameters, allowing for adaptive learning even in environments with uncertain parameters. Recent works highlight, however, that treating quantum noise and data uncertainty independently is inadequate since their compounded effect can severely diminish model performance. This has led to increased interest in integrated approaches that represent these uncertainties together. Methods including Monte Carlo sampling, Variational inference and ensemble-based methods are gradually modified to operate in quantum environments for better robustness and predictive confidence.

Overall, the literature shows a fundamental shift from purely performance oriented QML models to reliability orientated framework designs Capable of awareness of uncertainty and noise. It is this shift that is key to bringing tangible real-world applications of QML. The key idea of uncertainty-aware Quantum Machine Learning is that combining quantum computing with probabilistic methods not only increases the stability of a model, but also allows for interpretability by helping to make trustworthy decisions.

Study of Noise and Uncertainty in Quantum Machine Learning

Quantum Machine Learning (QML) systems are executed at the junction of quantum computing hardware and classical data-driven algorithms. Consequently, they are subjected to different sources of noise and uncertainty which together affect their reliability and performance. These sources must be understood to form strong uncertainty-aware QML models. It is useful to divide the noise and uncertainty in QML broadly into quantum hardware noise, classical data noise, and model-based uncertainty.

A. Quantum Hardware Noise

The most prominent and hardest challenge in QML is hardware noise, which comes because qubits are very delicate. State-of-the-art Noisy Intermediate-Scale Quantum (NISQ) devices have qubits that are very sensitive to environmental interactions, causing errors in computation. This noise has a lot of sources, one of them being DE coherence, during which quantum states lose their coherence when interacting with the environment. This leads to the decay of quantum information prior to computation having been performed.

In addition to DE coherence, an important source of error is gate error which arises from quantum logic gates not executing their operations with perfect fidelity. This accumulation of even small gate errors in parameterized circuits can distort results significantly since many QML models are based on sequences of quantum gates. On top of that, there is measurement noise since measuring a quantum state disturbs it by introducing stochastic spreads in the output. When combined, these factors generate a non-deterministic computational setting with sketchy reproducibility.

B. Classical Data Noise

Aside from hardware problems, QML systems are also limited by noise observed in the classical datasets used for training and evaluation. This sort of noise can be created from with errors in data collection, such as missing values, sensor inaccuracies or just architectural stochasticity. This uncertainty is commonly known as aleatoric uncertainty, because it comes from the data itself and cannot be eliminated completely.

The learning process of QML models can also be hindered by classical data noise. As quantum circuits are used to map classical data onto a quantum state, any errors in the input data will simply translate into the quantum computation. This can lead to wrong representations of features and in turn affect the performance of the model. When it comes to applications such as healthcare or financial forecasting, where data quality is of utmost importance even a small amount of noise can result in false positives.

C. Model Uncertainty

Model uncertainty or epistemic uncertainty is a more nuanced type of complexity that emanates from manifold factors, among which the learning model itself is one of limitations. This includes not enough training data, over fitting, under fitting and wrong assumptions about the underlying data distribution. Due to the rich structure of quantum circuits and difficulties in optimizing their parameters, model uncertainty in QML is exacerbated.

Here, the Variational quantum models are particularly prone to problems such as barren plateaus—optimizing over a region of landscape where the gradients vanish making training very hard. This results in an increased difficulty with parameter estimation and a loss of generalizability of the model to unseen data. This would imply that predictions given by QML models might be considered unreliable, particularly in case of few or noisy training data.

D. Interaction of Noise Sources

Interaction of the two distinct noise and uncertainty types is one of the most important characteristics of QML systems. Noise from quantum hardware and noise associated with classical data do not operate independently, but rather interact in a manner that actually magnifies total uncertainty in the overall system. As an example, noisy input data encoded into quantum states gets corrupted further by DE coherence and gate errors during computation. It makes training processes unstable, lowers accuracy and increases prediction variability due to this compounded effect.

Table 1: Sources of Noise and Uncertainty in QML Adapted from Sauerwein and De los Santos

Category	Source	Description	Impact on QML Systems
Quantum Hardware Noise	DE coherence	Physical processes introducing DE coherence to the quantum state	Reduces computation fidelity
	Gate Errors	Imperfect implementation of quantum operations	Accumulates errors in circuits
	Measurement Noise	Variability during quantum state observation	Produces inconsistent outputs
Classical Data Noise	Incomplete Data	Missing or insufficient data points	Leads to biased learning
	Measurement Errors	Inaccurate data collection	Distorts input features
	Data Variability	Natural randomness in real-world data	Introduces uncertainty in predictions
Model Uncertainty	Limited Training Data	Insufficient data for model learning	Poor generalization

	Model Assumptions	Incorrect assumptions about data distribution	Reduces accuracy
	Optimization Challenges	Lack of proper training plateaus (e.g., barren plateaus)	Slows convergence and reduces reliability

To summarize, QML is noisy and imperfect for many reasons including a combination of aspects across quantum hardware, classical data as well as the model it. This is a problem that must be tackled holistically, taking into account all sources at once. This can be seen as a way of going up the uncertainty ladder: Despite not being able to know what is causing specific errors, by identifying their origins and how they are interacting with one another, researchers could build more definitive QML systems that take uncertainties into account and would be able to provide reliable performance in noisy scenarios.

Uncertainty Quantification Techniques

UQ has a key role in improving confidence and interpretability of Quantum Machine Learning (QML) systems. In contrast to classical deterministic models, uncertainty-aware QML frameworks target producing probabilistic outputs that represent not just predictions per se but the level of confidence accompanying the prediction. In particular, this is crucial for noisy quantum environments, where the uncertainty due to both hardware imperfections and data variability can be large. QML models can lead to better and more reliable decision-making processes when paired with sound uncertainty quantification methods.

A. Bayesian Inference

Bayesian inference is one of the most common methods for uncertainty quantification. This approach also offers a principled framework for uncertainty modelling as instead of fixed values, model parameters are viewed as distributions over possible values. From the QML perspective, Bayesian inference combines prior knowledge or information with observed data to get an updated belief about model parameters.

Bayesian methods can be harnessed in quantum models (e.g., parameterized quantum circuits) so that uncertainty in circuit parameters can also be estimated. This creates predictive distributions, rather than single point predictions, allowing it to properly capture epistemic uncertainty. Bayesian methods are particularly useful when limited data is available, as they provide a formal approach to taking prior information into account and improving generalization.

B. Monte Carlo Sampling

One other great approach for uncertainty estimation is the Monte Carlo method. That is, through drawing many samples from some probability distribution and responding with those samples to approximate statistical properties like mean and variance. In QML context, Monte Carlo methods are mostly used for simulating quantum measurements and estimating expectation values with uncertainties.

In quantum circuits, repeated sampling can lead to a distribution of outcomes instead of a binary deterministic result. This way we can capture regarding both, quantum noise and uncertainties in the model. Monte Carlo (MC) sampling is another widely used method in Bayesian QML models for approximating posterior distributions when analytical solutions are intractable. This approach is computationally demanding but also highly general and accurate for uncertainty quantification.

C. Variational Inference

Variational inference is an optimization-based method to find an approximation to a complex probability distribution. The general gist behind Variational inference is to approximate the true posterior distribution (which is typically intractable, e.g. directly train a neural network that computes) with an alternative 'simpler' distribution (called the Variational approximation) and minimize a measure of divergence between both of them.

QML is especially advantageous in combination with Variational Quantum Algorithms (VQAs) that utilize parameterized quantum circuits [1]. It does so by allowing Variational parameters to represent uncertainty, thus making it possible to learn and estimate uncertainties efficiently even in high-dimensional quantum systems. A pragmatic solution for real world QML implementations is provided by Variational inference, which effectively balances computation and accuracy.

D. Confidence Interval Estimation

Interval Estimate confidence is a classic statistical method adapted to be implemented in QML systems. It is an expression of a value being predicted lying within a certain probability range. This is also a useful way of evaluating how much one can trust the outputs for applications.

As for quantum, outputs can be derived on the basis of running many measurements out of blocked quantum circuits. Because quantum measurements are fundamentally probabilistic, one must take multiple number of observations to obtain reasonable estimates for the expectation values. QML systems solve this problem in a

supervised way, given an answer q (e.g., whether the patient is sick or not) (Supervised Quantum Machine Learning | Qi skit), confidence intervals can be calculated so that QML systems express how certain they are about their answers, which is important to make decisions responsibly in high-level sectors such as healthcare and finance.

E. Integrated Uncertainty Modelling

Even these separate and complementary approaches to uncertainty quantification are often combined in modern QML systems. An example of this is the combination of Bayesian inference and Monte Carlo sampling to approximate posterior distributions, or Variational inference as a way to expand these techniques in larger systems. These hybrid strategies allow QML models to simultaneously characterize aleatoric and epistemic uncertainties.

Conclusion Uncertainty quantification methods are required for reliable QML systems that can work in a noisy world. Researchers are able to model uncertainty using techniques like Bayesian inference, Monte Carlo sampling, Variational inference and confidence interval estimation so that they not only predict the future but also express their certainty. This ability is crucial to trust-based decision-making for quantum applications in the wild.

Table 2: Uncertainty Quantification Techniques in QML

Technique	Core Principle	Application in QML	Advantages	Limitations
Bayesian Inference	Models uncertainty using probability distributions	Parameter estimation in quantum circuits	Captures epistemic uncertainty effectively	Computationally expensive for large models
Monte Carlo Sampling	Approximation using repeated random sampling	Simulation of quantum measurements and expectation values	Flexible and accurate estimation	High computational cost
Variational Inference	Approximates complex distributions via optimization	Integrated with Variational quantum circuits	Efficient for large-scale problems	Approximation may reduce accuracy
Confidence Interval Est.	Provides probabilistic range for predictions	Estimating reliability of quantum outputs	Easy interpretation	Requires multiple measurements

Rationale behind Noise-Aware Quantum Machine Learning Models

Noise-aware QML models are a significant step towards the use of quantum algorithms in real world settings, where noise and hardware imperfections can never be neglected. In contrast to idealized quantum models that assume perfect quantum operations, noise-aware methods consider effects due to DE coherence and gate errors as well as measurement uncertainty in the design and training of the model. Such a transition is crucial for the reliable performance in this NISQ era, where hardware constraints directly affect computational results. Noise-aware QML models include the Quantum Support Vector Machines (QSVMs), Quantum Neural Networks (QNNs), and various hybrid quantum-classical architectures, each with a distinct way of encoding noise with an aim to continuing learning.

Quantum support vector machines (QSVMs) build on and adapt classical support vector machines for the quantum world through quantum feature maps and kernel estimation methods. The classical data is mapped onto quantum states by embedding into a parameterized circuit and the data point similarity is calculated using quantum kernels in QSVMs. The kernel-based methods underlying QSVMs, as opposed to deep circuit architectures, make them ideal in noisy settings. Kernel estimation can be done with rather shallow quantum circuits so QSVMs are less sensitive to DE coherence and gate errors than deeper Variational models. Also, the noise-aware QSVM implementations include additional stabilizing techniques such as the use of a regularization strategy and evaluation of error-mitigated kernels which enhances classification accuracy. These properties favour QSVMs for near-term quantum devices, where a circuit depth has to be kept small in order to avoid noise accumulation.

A second prominent category of noise-aware QML models is Quantum Neural Networks (QNNs), which are analogous to classical neural networks but use parameterized quantum circuits. QNNs are composed of quantum gates organized into layers, with the parameters of the quantum gates being optimized during training to respond to tasks such as classification or regression, as well as generative modelling. Nonetheless, QNNs have a fundamental sensitivity to noise due to the many gate operations they rely on. In response, researchers have proposed noise-aware quantum neural networks (QNNs) with shallow circuit architecture designs, parameter sharing, and strong optimization techniques. As an example, shallow circuits prevent quibits from slow deciphering processes on time scales larger than circuit depth and a parameter-efficient design restricts the number of trainable variables leading the learnt ersatz to over fitting or being unstable. Furthermore, noise-injected optimization [] in which realistic noise models are injected during training also allows QNNs to learn fault-tolerant representations. All of these adaptations led to a great performance with QNNs even in noisy environments, thus rendering them as viable candidates for real-world QML applications.

Combining quantum circuits with classical machine learning components, hybrid quantum-classical models are an especially promising method for realizing noise-aware QML. These architectures generally have a quantum processor for transforming/encoding features and classical neural networks to perform optimization, make decisions, and generate outputs. Hybrid models take advantage of the strengths of quantum computing (e.g., representational completeness with high-dimensional features) while falling back on classical systems due to their robustness and scalability through division of labor. Instead, hybrid models are more noise-resilient by nature because classical components can override errors that the quantum circuits introduce. For example, classical post-processing layers may learn to clean the output of noisy quantum circuits, and mitigate hardware noise on a final prediction. Additionally, hybrid training frameworks provide feedback loops between quantum and classical elements facilitating dynamic updates to model parameters based on realized noise rates.

One of the key features in noise-aware QML models is that they allow to include error mitigation methods during learning. Techniques such as zero-noise extrapolation, probabilistic error cancellation and noise-adaptive loss functions are progressively being adopted in QSVMs [Cv07], QNNs [LL19] and hybrid architectures. These do not remove noise completely but again reduced its effect on model and help it to make better predictions. Ensemble methods, in which several quantum models are trained and then combined, have also been investigated to potentially average out errors caused by noise and increase robustness.

However, building data driven QML models that accurately capture noise scales up well with larger and more complex problems. The impact of noise is enhanced with a greater number of qubits as well as deeper circuits, effectively making contemporary mitigation techniques ineffective. Moreover, in the presence of noise, quantum models can be fragile and they should be designed to strike a delicate balance between expressiveness and stability: large (complex) models tend to perform very badly. Current work involves finding better designs for specific circuits, the use of adaptive noise models, and amplifying existing hybrid frameworks to alleviate some of these limitations.

SGC-22NOISE-AWARE MODELS: Qsvms, Qnns AND A WIDER LANDSCAPE⁴In this section we provide a brief overview of some important classes of noise-aware QML models. These models provide explicit methodology to remember noise and de-risking path, thus enhancing the robustness of QML systems. Noise-aware approaches will continue being necessary for scaling quantum machine learning applications in practical settings, especially as the quantum hardware evolves.

Hybrid Quantum–Classical Approaches

Introduction Hybrid quantum–classical approaches turned out to be among the most accessible and effective ways to deploy Quantum Machine Learning (QML) in practice. Such models combine quantum circuits with classical computing methods, which allows them to be trained to be both more efficient and robust against hardware constraints. Since current quantum systems are already NISQ devices, hybrid strategies offer a natural setting that efficiently combines both quantum and classical machinery.

A. Architecture of Hybrid Models

A typical hybrid model consists of two components: a quantum layer and a classical layer. The quantum part encodes input data into quantum states and transforms it using parameterized quantum circuits. These circuits take advantage of the quantum phenomena like superposition and entanglement to produce high-dimensional feature representations. While the classical component takes care of optimization, parameter updates and decision-making processes. This decoupling between the two roles allows for hybrid architectures to wear down some of the computational overheads on quantum hardware while preserving model expressiveness.

B. Training and Optimization Mechanisms

Hybrid models undergo an iterative loop between quantum and classical systems. A quantum circuit outputs results, depending on existing parameters and is evaluated with the use of a classical cost function. Gradient descent or evolutionary strategies are examples of classical optimization procedures based on the parameters updating for minimizing the loss. The feedback continues loop until the model converges. Classical optimizers are very practical on such tasks, since there is no efficient cost function that has to be evaluated when training quantum circuits (this might drastically overcomplicate the task due to various kinds of quantum noise and/or error correction required).

C. Robustness in Noisy Environments

Hybrid approaches are robust against noisy data. Specific portions of the model are actually used for quantum computations, meaning that imperfections in the hardware will generally have a much lower impact on the performance of the overall system. By learning a mapping that could improve prediction accuracy on the noise corruptions, classical layers can compensate for errors introduced in the quantum circuits. Hybrid models can also implement kinds of hybrid strategies, where realistic observation noise is prepared during optimization. This allows the model to tolerate imperfections and continue producing good results in a very noisy environment.

D. Applications and Practical Significance

Hybrid quantum-classical models have been utilized in multiple elements of classification, optimization, and generative modelling. They offer flexibility and scalability that makes them an excellent platform for near-term quantum applications, especially in some sectors where there are strict reliability requirements. While quantum hardware is still in the early development phase, hybrid strategies are anticipated to be an intermediary stage between noisy intermediate-scale quantum (NISQ) devices and large-scale fault-tolerant systems.

To sum up, hybrid quantum-classical methods offer a realistic route to QML implementation at scale. These models enable learning and decision making in noisy environments with good reliability, combining quantum computational advantages and classical stability.




Figure 2: Hybrid Quantum-Classical Architecture Overview

Error Mitigation Strategies

Error mitigation is an integral part of the Quantum Machine Learning (QML) and especially in this Noisy Intermediate-Scale Quantum (NISQ) era, where quantum hardware is by nature error-prone. In contrast, fundamental quantum error correction necessitates sophisticated fault-tolerant systems, whereas error mitigation methods seek to minimize the adverse influence of noise without significantly exacerbating hardware complexity. Such strategies improve reliability and performance of QML models so that they can operate effectively in noisy surroundings.

A. Zero-Noise Extrapolation (ZNE)

- The technique known as zero-noise extrapolation is one of the most common methods for suppressing quantum errors.
- One can think of it as artificially injecting the noise in a quantum circuit and expanding backward results to an ideal zero-noise case.
- This is generally done by stretching gate times or inserting more noisy operations to give different noise levels.
- A mathematical model estimates the noiseless solution by observing how outputs vary as noise increases.
- ZNE is therefore more compatible with near term quantum devices, as it does not increase qubit count.
- However, this does incur an increase in computational overhead with repeated executions of the circuit and can also be sensitive to assumptions made about the model.

B. Probabilistic Error Cancellation (PEC)

- Probabilistic error cancellation utilizes a probabilistic framework to conciliate noise with mathematical reversibility.
- The algorithm models noise as a quantum channel and uses inverse operations to mitigate these effects.
- This technique relies on a thorough description of the noise features in the quantum system.
- In theory, PEC can yield perfectly accurate output similar to ideal quantum computation.
- But it has a high computational cost because a large number of circuit samples are needed.
- This approach also incurs sampling overhead, which can scale exponentially as the circuit grows in size.

C. Noise-Aware Circuit Design

- Noise-aware circuit design (NACD) is all about designing quantum circuits that are intrinsically resistant to noise.
- This includes reducing the exposure to DE coherence and gate errors through minimizing circuit depth.
- These shallow circuits and optimised gate sequences maintain fidelity of the quantum states.
- These include qubit mapping and hardware-efficient ansatz to organize circuits according to the topology of quantum devices.

- Customizing circuit design to accommodate hardware limitations decreases the necessity of error correction after post-processing.
- This method is especially powerful when coupled with hybrid quantum-classical optimization methods.

D. Machine Learning-Based Error Mitigation

- Machine learning approaches are becoming increasingly popular to mitigate errors in QML systems.
- Models can be trained to recognize noise and output error-free predictions from noisy quantum data.
- Supervised learning techniques build correction models based on labelled datasets of noisy and ideal outputs.
- This enables Quantum circuits to be built which focus on selecting sequences of gates which minimize noise.
- This property makes these approaches very adaptable to changing noise scenarios.
- Integrating Bayesian inference also enables probabilistic modelling of noise, as well as generating uncertainties and robustness.

E. Hybrid Error Mitigation Strategies

- Indeed, in practice, implementing several error mitigations together tends to work even better than any single one.
- Zero-noise extrapolation, for instance can be combined with noise-aware circuit design to achieve accuracy whilst keeping a tab on computing expenses.
- Hybrid strategies are systems which have both hardware-level and algorithm-level optimizations.
- In particular, classical post-processing methods can improve the quantum outputs, making an integrated system more reliable.

Error mitigation strategies are key to narrowing the gap between theoretical quantum advantage and practical realizations. Zero-noise extrapolation, probabilistic error cancellation and noise-aware circuit construction are just a few strategies that can also be applied to overcome hardware imperfections. These techniques provide adaptive and scalable solutions to enhance the performance of QML systems when used along with machine learning based approaches. This calls for further research in quantum machine learning and error mitigation to boost operational confidence and robustness with quantum technology moving forward in real world applications.

Uncertainty-Aware Optimization

Adaptive optimization under uncertainty is the key ingredient of Quantum Machine Learning (QML) in practice, applicable when quantum noise and data variability both pose challenges to performance. Standard optimization methods assume a deterministic setting, an idealized circumstance that is far from the reality of quantum systems. This limitation is temporal given that uncertainty-aware optimization directly includes the uncertainty information during the optimization, which can allow models to consider various other factors and finally make more reliable robust decisions⁸. Such techniques are required to stabilize training and improve predictive performance on Noisy Intermediate-Scale Quantum (NISQ) devices.

A. Robust Optimization

Robust optimization designs models so that they are resilient to worst-case scenarios of uncertainty. Rather than optimizing for a single expected outcome, this method encompasses wide variations in input data and system parameters. The robustness optimization of QML guarantees that quantum models can keep efficacy under circumstances that they get noise, perturbations or adversarial conditions.

It is useful for working with aleatoric and epistemic uncertainties. For instance, in the context of a quantum classification task, robust optimization regulates some model parameters to lessen the effect of noisy inputs and uncertain measurements. The model can still differ but less affected by disturbance and more stable during the inference stage by modelling variability during train. Strong optimization tends to provide conservative solutions because it focuses on reliability not peak performance.

B. Risk-Sensitive Learning

Risk-sensitive learning adds a probabilistic aspect to the optimization process by weighting consequences according to their risks. Traditional optimization targets average performance, whereas risk-sensitive methods target minimization of losses in high-risk situations. This is especially critical in such applications as healthcare, finance, and autonomous systems where incorrect predictions can cost lives.

For example, risk-sensitive learning can be applied in QML through changes to the loss function that place greater penalties on uncertain or high-risk predictions. For example, high variance or low confidence predictions may result in larger penalties leading the model to stick towards safer and more reliable outputs. This works in a similar fashion to uncertainty quantification methods (Bayesian inference for example) that characterize the

uncertainties with probabilistic distributions. The optimization of QML models encompasses a balance between performance and reliability using risk-aware methods.

C. Distribution ally Robust Optimization)

Distribution ally Robust Optimization (DRO) is a generalization of simple robust optimization, in which there is uncertainty not only in the data but also in its probability distribution. Instead of presuming a static data distribution, DRO method seeks to improve the model performance over a collection of potential distributions in some bounded uncertainty area. In fact, this is especially important in QML since classical data noise and/or quantum measurement variations can change the effective distribution of data [8].

DRO gives strong performance guarantees for a model under distributional shifts, making it very practical for real-world scenarios where the data condition can change over time. In quantum settings, we can use DRO to optimize Variational quantum circuits under combined uncertainty in coded input as well as measurement results. More complicated to compute than normal optimization methods, but it provides better generalization and robustness.

D. Integration with Quantum Models

In QML, classical optimization framework is adapted to quantum circuits leveraging uncertainty-aware optimization techniques. This integration is enabled by hybrid quantum-classical models, which allow classical optimizers to take indirect uncertainty measures under the interaction with quantum components. Uncertainty-aware optimization in QML is typically implemented using techniques such as stochastic gradient methods, adaptive learning rates and probabilistic loss functions.

Combining these optimization strategies with error mitigation and uncertainty quantification techniques improves overall system performance. For instance, noise-aware training can be used to enhance the stability of robust optimization approaches, and confidence estimates from a probabilistic model can help augment risk-sensitive learning.

All in all, uncertainty-aware optimization is a strong framework to enhance the reliability and robustness of QML systems. Robust optimization, risk-sensitive learning and distribution ally robust optimization are some well-studied techniques that facilitate model performance in uncertain and noisy conditions. These approaches provide the means to embed uncertainty directly into the optimization process, thereby allowing QML models to yield reliable, repeatable and accurate predictions — integral requirements for real-world deployment.

Explainability On Top Of Uncertainty-Aware Quantum Machine Learning

Explainability is a key prerequisite for the uptake of Quantum Machine Learning (QML) systems, especially in applications where the decisions have major real-world ramifications. But as the QML models started to use more complicated elements such as a composition of quantum circuits and probabilistic learning mechanisms, explaining how it arrives at a certain prediction becomes both challenging and crucial. In the case of uncertainty-aware QML, Explainability extends beyond interpreting predictions to understanding and communicating the uncertainty associated with predictions. Having both is essential to foster trust, provide transparency and allow for informed decisions in noisy environments.

Common techniques include feature importance, saliency maps and model-agnostic interpretation methods for traditional machine learning. The inherent properties of quantum systems such as superposition and entanglement apply additional complexity when these concepts are extended to QML. Such properties render intermediate states in quantum circuits hard to observe or interpret. Consequently, interpretability in QML is mainly indirect, such as probing the measurement results, investigating sensitivities of parameters and distributions of outputs.

Interpretability of probabilistic predictions is at the centre of Explainability in uncertainty-aware QML. Instead of a single point, uncertainty-aware systems return an output in the form of a probability distribution or confidence interval (as this is what we know about our model). These outputs can serve as indicators of how trustworthy predictions are, allowing users to make informed decisions based on the risk profile of certain options. Bayesian inference is one technique that lies at the heart of this broad idea, because it allows to include uncertainty in model parameters and predictions. The model confidence communicated through posterior distributions can allow users to better understand when predictions from the model may not be reliable.

An additional facet of Explainability is the certification of quantum circuit behavior. Parameterized quantum circuits serve as the basic computational units, found in QML models, where they convert input data into output predictions. Analyzing the contribution of single gates, qubit interactions & parameter value is crucial to understand how these circuits process information. Analysis of parameter sensitivity or gradient based methods can identify

which portions of the circuit have more influence on model performance. This information can be exploited to simplify circuits, robustify them and better interpret them.

Hybrid quantum-classical models provide further avenues for interpretability through the use of classical components which are inherently more interpretable. For instance, means to achieve inner interpretability through classical neural networks or decision trees incorporated within hybrid architectures may elucidate feature importance's and discrete decision boundaries while supplementing the functionality of the quantum component. With this, researchers will get the opportunity to connect quantum complexity and classical interpretability in QML systems, which should be user-friendly.

Explainability is also important for debugging and refining QML models. Prioritizing entropy sources and their effects on predictions can inform how to better tune model architectures or training methods for that problem space. If a model often shows high uncertainty in certain areas of the input space, this can be due to lack of adequate data for training or excessive noise at that time step of the quantum circuit. By tackling these challenges, better and more robust models can be constructed.

Additionally, Explainability is crucial for compliance with ethical and regulatory standards in high-stakes areas like healthcare, finance, and autonomous systems. In these domains, where uncertainty is part of the equation, decision-making systems must be clear on how they arrive at their outputs. For such applications, it is important to have QML models that are transparent in terms of communicating not only predictions but also the associated confidence levels to enable stakeholders in assessing risks and making decisions with a level of assurance.

We end with conclusions which describe why Explainability is a rich concept in uncertainty-aware QML-predictions, quantum circuit behaviour or the related uncertainty estimates. Bringing together probabilistic reasoning, sensitivity analysis, as well as hybrid modelling methods will support making QML systems more understandable and reliable. Lesson 2: With the rapid evolution of this field in recent years, more sophisticated and scalable Explainability approaches need to be developed so that quantum machine learning can immediately be deployed responsibly and successfully in our real-world environments.




Figure 3: Explainability on Top of Uncertainty-Aware Quantum Machine Learning

This Article Is a Part of Our Quantum Computing For Beginners Series and Are Based On Uncertainty-Aware Quantum Machine Learning

Uncertainty-aware quantum Machine Learning (QML) is poised to impact multiple high-stakes sectors by allowing assurance in decision making under the presence of noise and uncertainty. Integrating quantum computational advantages with probabilistic modelling, these systems facilitate not just accurate predictions but also confidence estimates, which are indispensable in many high-stakes real-world applications. This provides QML advantages in any environment with missing data, noisy data, or rapidly shifting data, as uncertainty is explicitly modelled. Some of its key application areas are seen as being in healthcare diagnostic, financial risk assessment, cybersecurity, and autonomous systems.

A. Healthcare Diagnostics

Healthcare is an area where making decisions to do feature can be a matter of life and death, as a result Precision is of utmost importance. For instance, in the field of medical image analysis and disease prediction/individualized treatment planning, there is significant potential for using uncertainty-aware QML models. By estimating uncertainty they also begin to show MOOCs the confidence with which diagnoses were made, providing clinicians with supplementary information needed for better-informed decision making. For instance, if a QML model was trained with medical imaging data and identifies certain areas where it did not have enough

certainty about the measurements, these can then be flagged for immediate investigation or further testing. This minimizes the chances of misdiagnosis and enhances patient outcomes. In addition, such quantum-enhanced models can be trained on more complex biological datasets at a faster pace and with greater accuracy.

B. Financial Risk Assessment

The financial industry is a particularly uncertain and dynamic environment, where having an accurate risk assessment in place can make or break profitability. Portfolio optimization, fraud detection and prediction market can be done using uncertainty-aware QML. These systems can create probabilistic forecasts, as they are not deterred by market volatility or result of various external factors that affect financial data. This enables financial institutions to make risk-sensitive decisions like changing investment strategies or detecting losses. Moreover, unlike classical approaches, the models of QML can make process the large-scale financial datasets with greater efficiency, providing hidden patterns and correlations that cannot be easily accessed by classical methods.

C. Cybersecurity Threat Detection

Uncertainty-aware QML can have a great role in another area, cybersecurity. Much of the data behind detecting cyber threats is noisy and incomplete—network traffic perhaps or user behavior. But these subtle patterns and anomalies are often difficult to find without sophisticated tools — tools like QML models that can examine complex datasets and surfaces hidden insights that ultimately enhance threat detection. With the addition of uncertainty estimates, these systems can better separate threats from non-threats. In this case, a model may raise an alert for a potential intrusion with say, 70% confidence and in turn, risk-based security analysts can prioritize responses directly. This enhances both the speed and accuracy of cybersecurity operations.

Table 3: Applications of Uncertainty-Aware QML

Application Domain	Use Case	Role of Uncertainty-Aware QML	Benefits
Healthcare Diagnostics	Disease detection, medical imaging	Provides confidence estimates for diagnoses	Reduces misdiagnosis and improves patient safety
Financial Risk Assessment	Portfolio optimization, fraud detection	Models market uncertainty and financial risk	Enables better investment decisions and reduces losses
Cybersecurity	Intrusion detection, anomaly detection	Identifies threats with probabilistic confidence	Reduces false positives and improves response time
Autonomous Systems	Self-driving vehicles, robotics	Supports decision-making under uncertainty with risk estimation	Enhances safety and operational reliability

Experimental Design And Evaluation

A rigorous experimental design that distinguishes between limitations imposed by quantum hardware, as well as uncertainties in classical data, is fundamental for the development and validation of uncertainty-aware QML models. In contrast to traditional machine learning experiments, assessing an application of QML requires noticing the stochastic character of quantum measurements, hardware noise, and hybrid quantum-classical workflow integration. Thus, a detailed experimental system is necessary to obtain reliable, reproducible, and meaningful results in practical applications.

Dataset and problem domain selection is the first step in experimental design. In uncertainty-aware QML, datasets should be crafted to mirror the bottlenecks of real life that includes noise, missing and variable values. Due to the ability to introduce different levels and types of uncertainty in a controlled way, synthetic datasets are often used for controlled experiments¹⁰³. But the other side of this coin in real world data sets, which is equally important for investigating practical application considerations, especially in domains such as healthcare, finance and cybersecurity. Dataset choice directly affects the evaluation of both prediction accuracy and uncertainty quantification.

Choosing quantum models and circuit architectures is another key element of experimental design. They are common in Variational quantum circuits, and hybrid quantum-classical models since they run very well on the current state of NISQ devices. Similarly, the depth and structure of quantum circuits must be designed to provide an ideal trade are between expressiveness and noise resilience. This is particularly beneficial for shallow circuits, as in noisy environments any accumulated error from DE coherence and imperfect gate operations will be present. Furthermore, initialization and optimization strategies are important for model performance because of problems such as barren plateaus.

Uncertainty-aware QML relies on evaluation metrics that go beyond traditional ones such as accuracy and precision. These metrics keep being relevant to evaluate predictive performance; however additional criteria are needed that enable assessing uncertainty estimation and robustness. Calibration metrics, including expected

calibration error (ECE), evaluate the degree to which predicted probabilities follow true outcomes. Robustness Metrics: These metrics evaluate the stability with which a model is replicating its performance across different amounts of noise and perturbations in data. Moreover, prediction entropy and variance deliver information about the amount of uncertainty in model predictions. This is important for assessing whether a model can give sensible confidence estimates.

Additionally, the experimental framework should simulate quantum hardware by implementing realistic noise modelling. The noise models commonly include DE coherence, gate errors and measurement noise typically present in NISQ devices. This way, researchers can find out how well their models function within practical restrictions by including these disturbance designs into reproductions or using genuine quantum equipment. Experimental error mitigation techniques such as zero-noise extrapolation and probabilistic error cancellation are typically used to obtain more accurate results with the goal of predicting their efficacy.

Another important aspect in QML experiments is Reproducibility. Quantum measurements have stochastic nature; hence performing several repeats of the same experiment might give us different results. This means that, to mitigate this, we usually run experiments multiple times and present performance as statistics, e.g. average values. Additionally, inference through confidence intervals and variance analysis gives an illustration of the degree by which the outcomes would be firm and reliable. This technique guarantees that we are making confident conclusions based on our experimental results and that they are not the result of random variability.

When benchmarking uncertainty-aware QML models, a comparative analysis is usually made with classical machine learning approaches. Involves testing both kinds of models on the same datasets and under similar conditions, to compare them directly. These comparisons shed light on the strengths and weaknesses of QML, specifically in terms of noise and uncertainty. Hybrid models tend to exhibit greater robustness in a lot of circumstances; purely quantum models can be beneficial for certain computation tasks.

Lastly, visualization and interpretability become essential for evaluating an experiment. Visualizations related to uncertainty distributions, error rates, and performance trends are helpful for researchers in understanding the behavior of models. They are especially useful in finding patterns and diagnosing problems, as well as communicating results effectively.

Collectively, these four key features of experimental design and evaluation in uncertainty-aware QML are part of a holistic approach that includes careful selection of a dataset, sensible design of the model, and realistic simulation of noise during training, and improved metrics more sensitive to both epistemic and aleatoric uncertainty. In conclusion, with proper handling of all these points, one can create and apply models in noisy settings that are accurate but also reliable and robust. Such a holistic metric assessment framework is fundamental to moving QML from the theoretical into pragmatic applications in practice.

Limitations and Challenges of Uncertainty-Aware QML

While the incorporation of uncertainty in Quantum Machine Learning (QML) shows great promise, there are various challenges and limitations that must be overcome for QML to achieve wide-scale adoption and deployment. Both quantum hardware limitations and the inherent complexity of incorporating uncertainty into any system drive these challenges. Solving these problems is necessary for the progress of QML from fundamental research to practical applications.

The major hurdle is the performance of existing quantum hardware, especially Noisy Intermediate-Scale Quantum (NISQ). Current quantum hardware consists of a limited number of qubits with low cycling times and high error rates. These limits impose a smaller and shallower region of quantum circuits possible to be executed, severely limiting the complexity that a QML model can achieve. This means that even those algorithms which are conceptually very powerful cannot run properly on existing machines just yet. Moreover, both noise and DE coherence lead to variability in the quantum computations that precludes obtaining the same results repeatedly [117].

A second big limitation is difficult to model and measure uncertainty in quantum systems. To begin with, classical machine learning already has well-defined methodologies for uncertainty quantification — such as Bayesian approaches and probabilistic modelling — but extending these techniques to quantum environments is by no means trivial. Quantum systems themselves are probabilistic in nature, and this can be hard to tell apart from uncertainty due to noise or model miss-specification. This overlap makes interpreting results more difficult, and separating meaningful uncertainty from random noise is often impossible.

Another important challenge in uncertainty-aware QML is scalability. The complexity of computation and uncertainty modelling increases exponentially with the size of quantum systems. Training quantum models at scale incurs a large computational cost, especially with approaches like Monte Carlo sampling or Variational inference.

Meanwhile, hybrid quantum-classical models also increase the overheads associated with iterative communication between quantum and classical components, although they are more practical. This could not only escalate the latency but can be less efficient for real-time applications.

Hence we could say that QML models suffer from more constraints because they also deal with optimization challenges. Word on the street is: Variational quantum algorithms usually used in QML suffers from barren plateaus (i.e., regions where gradients go to zero, which makes training extremely difficult even when one lacks huge amount of classical computation resources) In environments with noise, this issue escalates $\Rightarrow$ uncertainty about gradient estimation can cause convergence to be either unstable or sluggish. Consequently, optimizing parameters for use in quantum circuits is among the most challenging tasks, especially as models become more complicated.

A third, crucial challenge is the absence of standardized benchmarks and evaluation frameworks for uncertainty-aware QML. At each stage, classical ML has clearly delineated datasets and performance metrics while the QML field is still developing its evaluation standards. Sources of variability include differences in hardware platforms, noise models and experimental setups that make it cumbersome to directly compare results across studies. This inconsistency impedes progress and sows doubt over how much promise these QML approaches really offer.

There are also important limitations with respect to interpretability and Explainability. As mentioned in the earlier sections, it is fundamentally hard to understand decision-making by QML models—largely due to complexity backed behind quantum computations. This complexity further increases when uncertainty is incorporated into the model.

References

- [1] Nielsen, M.A. and Chuang, I.L. (2010) *Quantum Computation and Quantum Information*. Cambridge: Cambridge University Press.
- [2] Seth Lloyd (2013) *Quantum machine learning*. ArXiv preprint arXiv: 1307.0411.
- [3] Jacob Belmonte et al. (2017) *Quantum machine learning*. Nature, 549(7671), pp.195–202.
- [4] Maria Schulz and Kilogram, N. (2019) *Quantum machine learning in feature Hilbert spaces*. Physical Review Letters, 122(4).
- [5] Michele Carazo et al. (2021) *Variational quantum algorithms*. Nature Reviews Physics, 3, pp.625–644.
- [6] Rotech Havelock et al. (2019) *Supervised learning with quantum-enhanced feature spaces*. Nature, 567, pp.209–212.
- [7] Edward Farsi et al. (2014) *A quantum approximate optimization algorithm*. ArXiv preprint arXiv: 1411.4028.
- [8] Nathan Weise et al. (2014) *Quantum deep learning*. ArXiv preprint arXiv: 1412.3489.
- [9] Patrick Rebentrost et al. (2014) *Quantum support vector machine for big data classification*. Physical Review Letters, 113(13).
- [10] Marcello Benedetti et al. (2019) *Parameterized quantum circuits as machine learning models*. Quantum Science and Technology, 4(4).
- [11] Kosice Mitral et al. (2018) *Quantum circuit learning*. Physical Review A, 98(3).
- [12] Jarrod R. McLean et al. (2018) *Barren plateaus in quantum neural network training landscapes*. Nature Communications, 9(1).
- [13] Sugar Endo et al. (2018) *Practical quantum error mitigation for near-future applications*. Physical Review X, 8(3).
- [14] Kristen Temme et al. (2017) *Error mitigation for short-depth quantum circuits*. Physical Review Letters, 119(18).
- [15] Zhukov, S. et al. (2020) *Noise in quantum computing*. Journal of Physics A.
- [16] Murphy, K.P. (2012) *Machine Learning: A Probabilistic Perspective*. MIT Press.
- [17] Bishop, C.M. (2006) *Pattern Recognition and Machine Learning*. Springer.
- [18] Kevin P. Murphy (2022) *Probabilistic Machine Learning*. MIT Press.
- [19] Goodfellow, I., Bagnio, Y. and Carville, A. (2016) *Deep Learning*. MIT Press.
- [20] Eye Huller Meier and Aegean, W. (2021) *Aleatoric and epistemic uncertainty in ML*. Machine Learning Journal.
- [21] Gal, Y. (2016) *Uncertainty in Deep Learning*. PhD Thesis, University of Cambridge.
- [22] Carl Edward Rasmussen and Williams, C.K.I. (2006) *Gaussian Processes for Machine Learning*. MIT Press.
- [23] IBM Quantum (2025) *Quantum Machine learning advancements*.
- [24] Google Quantum AI (2024) *Quantum AI research overview*.
- [25] Microsoft Quantum (2025) *Quantum computing frameworks*.
- [26] Regatta computing (2024) *Hybrid quantum computing platforms*.
- [27] Landau (2024) *Photonic quantum machine learning*.
- [28] Scott Aaronson (2013) *Quantum computing since Democritus*. Cambridge University Press.
- [29] John Atrous (2018) *the Theory of Quantum Information*. Cambridge University Press.
- [30] Daniel Gateman (2010) *an introduction to quantum error correction*.
- [31] Peter Witter (2014) *Quantum Machine Learning: What Quantum Computing Means to Data Mining*. Academic Press.
- [32] Aram Harrow et al. (2009) *Quantum algorithm for linear systems of equations*. Physical Review Letters.
- [33] Love Grover (1996) *a fast quantum mechanical algorithm for database search*.
- [34] Peter Short (1994) *Algorithms for quantum computation: Discrete logarithms and factoring*.