

Journal Pre-proof

Quantum artificial intelligence for mission-critical systems: Foundations, architectural elements, and future directions

Siva Sai, Rajkumar Buyya



PII: S0167-739X(26)00236-0

DOI: <https://doi.org/10.1016/j.future.2026.108602>

Reference: FUTURE 108602

To appear in: *Future Generation Computer Systems*

Received date: 12 November 2025

Revised date: 10 May 2026

Accepted date: 12 May 2026

Please cite this article as: S. Sai and R. Buyya, Quantum artificial intelligence for mission-critical systems: Foundations, architectural elements, and future directions, *Future Generation Computer Systems* (2026), doi: <https://doi.org/10.1016/j.future.2026.108602>.

This is a PDF of an article that has undergone enhancements after acceptance, such as the addition of a cover page and metadata, and formatting for readability. This version will undergo additional copyediting, typesetting and review before it is published in its final form. As such, this version is no longer the Accepted Manuscript, but it is not yet the definitive Version of Record; we are providing this early version to give early visibility of the article. Please note that Elsevier's sharing policy for the Published Journal Article applies to this version, see: <https://www.elsevier.com/about/policies-and-standards/sharing#4-published-journal-article>. Please also note that, during the production process, errors may be discovered which could affect the content, and all legal disclaimers that apply to the journal pertain.

© 2026 Published by Elsevier B.V.

Quantum Artificial Intelligence for Mission-Critical Systems: Foundations, Architectural Elements, and Future Directions

Siva Sai, Rajkumar Buyya *Fellow, ACM; Fellow, IEEE*

Abstract—Mission critical (MC) applications such as defense operations, energy management, cybersecurity, and aerospace control require reliable, deterministic, and low-latency decision making under uncertainty. Although the classical Artificial Intelligence (AI) approaches are effective, they often struggle to meet the stringent constraints of robustness, timing, explainability, and safety in the MC domains. Quantum Artificial Intelligence (QAI), the fusion of artificial intelligence and quantum computing (QC), can potentially provide transformative solutions to the challenges faced by classical ML models. QAI is a broader umbrella than Quantum Machine Learning (QML) and additionally includes quantum optimization, search, and reasoning; we use QAI throughout the paper for the field at large, and QML only for learning-specific subroutines. The principal contributions of this work are: (i) a systematic survey of QAI methods analyzed through the lens of MC requirements like certification, robustness, and timing; (ii) a conceptual quantum cloud resource management and scheduling framework with deployment assumptions, complexity analysis, and failure-mode discussion; and (iii) an identification of the gaps between current QAI capabilities and MC systems requirements. We also propose a conceptual model for management of quantum resources and scheduling of applications driven by timeliness constraints. We discuss multiple challenges, including trainability limits, data access, and loading bottlenecks, verification of quantum components, and adversarial QAI. Finally, we outline future research directions toward achieving interpretable, scalable, and hardware-feasible QAI models for MC application deployment.

Index Terms—Quantum computing, Quantum artificial intelligence, Mission critical systems, Aerospace, Defense, Disaster management, Cybersecurity, Energy and grid management, Technical foundations

I. INTRODUCTION

Mission critical (MC) applications are those that must not fail because failure can lead to severe harm, injury, loss of life, large economic damage, or major outages [1]. An MC system aim to build a structured safety case that traces hazards, requirements, designs, and tests, so that acceptable risk is argued with evidence rather than based on assumption. An air traffic control system, where even a slight malfunction can endanger hundreds of lives is an example of an MC application. The industrial control systems in nuclear plants and power grids are also mission critical applications, since their failure could cause catastrophic accidents or large-scale blackouts. Machine Learning (ML) has become a default

approach for perception, decision support, and prediction, with its applications ranging from sensor monitoring to triaging incidents. Classical ML uses hand-crafted features with models like Support Vector Machines (SVMs) and gradient boosting, while Deep Learning (DL) learns features directly from the data. This capability of DL enabled breakthroughs in vision, speech, and control. Both ML and DL-based systems have a significant number of applications in MC systems, but these gains are not without the weaknesses that are critical in MC settings. The models might be overconfident in out-of-distribution (OOD) inputs, which is risky in early-warning systems and anomaly screens. Research [2] also showed that carefully altered adversarial data can fool the deep learning models, which is an integrity concern for autonomy stacks and safety monitors. The recent DL models are data-hungry and computationally heavy, and hence, acquiring labels for rare failures can be very expensive. Quantum Machine Learning (QML) models have the potential to address some of these limitations in MC applications, particularly for specific tasks where quantum feature spaces or optimization offer advantages over classical approaches.

Quantum Artificial Intelligence (QAI) is an umbrella term for the intersection of quantum computing and artificial intelligence. It is important to note that QAI is a much broader field than QML. QAI encompasses quantum optimization, quantum search, quantum reasoning, fuzzy quantum logic, among others [3], [4]. The scope of this paper is particularly limited to QAI methods applied to mission-critical systems, and we use QAI when referring to the field at large. QAI uses a quantum processor for particular steps in the learning process, like data preparation, measurement, and expectation estimation. The remaining parts of the hybrid pipelines, like data handling, logging, optimization, and control, remain classical. The key idea behind QAI is that by representing a classical input in quantum states, the device would be able to estimate overlaps that can act as useful kernels for standard ML classifiers. QAI leverages quantum phenomena like superposition and entanglement for processing vast amounts of information simultaneously, allowing better exploration of complex solution spaces and faster computation. MC tasks often face imbalanced, scarce labels (near-misses, rare faults) and are supposed to operate with tight assurance (bounded behavior, traceability). QAI can potentially help with these cases in concrete ways. QAI models can potentially act as powerful feature-map engines that enhance the separability of features even with limited

Siva Sai and Rajkumar Buyya are with the Quantum Cloud Computing and Distributed Systems (qCLOUDS) Laboratory, School of Computing and Information Systems, The University of Melbourne, Australia (e-mail: {sivasaireddy.naga, rbuyya}@unimelb.edu.au).

examples. With Quantum AI models, certain learning tasks require fewer experiments, particularly when the data comes from instruments (calibration, sensing, and characterization). Many mission critical tasks reduce to reconfiguration or hard scheduling with strict latency budgets. Variational quantum optimizers like Quantum Approximate Optimization Algorithm (QAOA) provide fixed-depth circuits that balance runtime and solution quality. When training in MC systems that must cross organizational boundaries, QAI can make use of QC techniques like quantum cryptography and blind quantum computing to protect queries, data, and model updates.

To the best of our knowledge, this work offers one of the early systematic syntheses of QAI methods for mission-critical systems. We also propose a resource management framework specifically in the context of mission-critical systems. The primary purpose of this work is to synthesize existing QAI for mission-critical systems research. In addition, we also offer a conceptual architectural contribution through the quantum resource management framework. The paper is conceptual and architectural in scope and we do not include original experiments or benchmarks; we have identified empirical validation on emerging quantum cloud platforms as a future direction. The rest of the paper is organized as follows: In Section II, we provide a brief background of quantum computing and mission critical systems. Section III elucidates quantum machine learning briefly from multiple perspectives. In Section IV, we provide core mechanisms and algorithmic principles of QAI in MC systems, followed by applications in Section V. Section VI provides an exploration of the positioning of QAI for MC systems in the industry in terms of deployment. In Section VII, we propose a framework to handle the unique properties and constraints of quantum computing in resource management and scheduling systems. We outline several challenges and future directions in Section VIII, and finally conclude the paper in Section IX.

II. BACKGROUND

A. Quantum Computing

Quantum computing deals with algorithms and systems that manipulate information in quantum states, which exploit two underlying principles of quantum physics - superposition and entanglement in order to transform the probability amplitudes in ways that differ from classical bit operations. Noisy Intermediate Scale Quantum (NISQ) view is a widely used perspective in today's quantum hardware, where devices with tens-hundreds of qubits run non-trivial circuits, but limited coherence and gate noise limit the reliability and circuit depth. This framing stresses the role of hybrid quantum-classical workflows and sets realistic expectations for near-term applications.

Quantum Models: There are two models that are widely used to think about quantum computers from application composition perspective: gate model and adiabatic/quantum annealing. The gate model builds a program from a small set of gates (primitive operations), analogous to logic gates in classical chips. DiVincenzo et al. [5] set out a classical

checklist of what hardware quantum hardware must provide - universal gates, scalable qubits, long stability, the ability to measure, and reliable ways to move information. In adiabatic or quantum annealing model, the problem is encoded in the energy landscape of a quantum system and guided towards a low-energy (good) solution. Research shows that this formulation is computationally equivalent to the gate model under ideal conditions [6].

Since the quantum states are fragile, large programs demand fault tolerance. In this connection, the threshold theorem states that if the individual operations are accurate enough, the error-correcting codes can suppress errors, thus allowing the running of very long computations also. Among many codes, the surface codes [7] are prominent as they work on 2-D grids with only local connections and have relatively high error thresholds. This makes them appealing for future scale systems.

Quantum Architectures: There are five leading approaches to building today's quantum hardware (refer Figure 1), and all of them have tradeoffs in terms of scale, control, and connectivity. Superconducting circuits [8] are based on microwave chips and enable fast gates and easy on-chip working, but have the issues of scaling and noise. Trapped ions [9] achieve high fidelity by controlling individual atoms with lasers. A drawback of this approach is that the larger systems can be complex and slow to route. Natural atoms/Rydberg arrays [10] achieve scalable layouts with strong and switchable interactions by arranging many atoms in optimized tweezers. Single photons and integrated optics are used in Photonic platforms [11], and they promise measurement-based logic and room-temperature operation. Semiconductor spin qubits [12] leverage chip fabrication and store information in electron/hole spins in silicon quantum dots, but they are still maturing in control at scale. Apart from the above solid-state and photonic platforms, chemical pathways in the liquid phase have been explored where the fuzzy logic is processed using thermalized quantum mixed states [13]. Although this unconventional architecture remains at an early stage, it broadens the landscape of physical substrates for quantum information processing.

Given the current limitations, the most practical quantum workflows are hybrid wherein the quantum processor executes a bounded task it is good at, like estimation overlaps or expectations, or preparing special states, while a classical processor handles the rest of the processes - data ingestion, optimization, and decision logic. Recent research reveals emphasis on the split in the tasks, the need to report noise levels, the number of measurements, and resource documents so that the results are reproducible and comparable across devices.

B. Mission Critical Systems

The mission critical systems are the ones that must not fail because failure can cause large economic losses, major service outages, and injury or death. A few clear examples include flight control, power-grid protection, and railway signaling. The required aspects of an MC system include reliability (keeps working), safety (no unacceptable harm), integrity (no

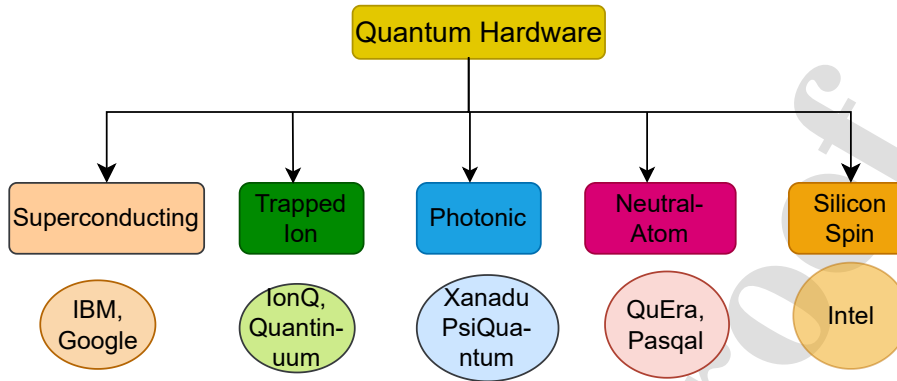


Figure 1: Leading approaches for quantum hardware realization and key industry players

improper changes), availability (ready when needed), and maintainability (can be fixed quickly) [1]. In the MC systems, faults (root causes) are differentiated from errors (incorrect internal state) and failures (outwardly visible problems).

Timing is a defining constraint in MC systems - the systems have to respond within guaranteed deadlines. Classical real-time scheduling theory allows assigning priorities so that the periodic tasks meet their deadlines on a processor, the theory of which is still used in automotive controllers, industrial automation, and avionics. In MC systems, the timing analyses are paired with monitoring and redundancy so that the system continues operating safely even if the components of the system misbehave. Various standards are set forth in terms of the mission critical systems to ensure the required aspects in the systems. IEC 61508 defines safety integration levels and a safety lifecycle that scale the development rigor and risk testing. ISO 26262 is adopted in automotive systems and defines automotive safety integrity levels with requirements on software, hardware, and validation. SO-178C standard specifies design assurance levels in aviation, which are linked to failure severity, and also mandates strong testability based on requirements through tests.

Across all MC domains, a common thread is that the hazard analysis must be done earlier, traceability from hazards must be maintained, an assurance level matching risk must be chosen, and investment in verification and validation must be made. So, for any new technology like quantum AI, acceptance in MC systems requires fitting into this framework. To integrate the QAI components into these certification workflows, they should be treated as software items subject to the same assurance level the system supports. Under IEC 61508, this will imply assigning a Safety Integrity Level (SIL) to the QAI component based on the hazard it mitigates and providing evidence of system capability (testing coverage, design reviews) at that level. The QAI component would require a Design Assurance Level (DAL) with traceability from requirements to test cases under DO-178C. It implies, in practice, that the quantum circuits be version controlled, i.e. their transpiled outputs be verified against reference implementations with equivalence checking

tools such as QCEC (Quantum Circuit Equivalence Checking) [14]. The stochastic nature of quantum measurement should also be bounded through repeated-shot statistical guarantees documented in the safety case.

III. QUANTUM MACHINE LEARNING AT A GLANCE

Quantum machine learning focuses on learning methodologies that use quantum information techniques at important steps like data representation, similarity measurement, or model optimization. Early reviews on QML helped in shaping the field, while cautioning against over-claiming. A well-cited early review on QML by Biamonte et al. [15] placed QML as its own field (in the year 2017) and noted that any practical gains would depend on proper assumptions on hardware and data access. In 2019, the work by Schuld et al. [16] revealed a central idea in the field - by preparing the classical data as quantum states, the quantum devices can realize kernels that are associated with high-dimensional feature spaces (which are very costly to evaluate classically). These kernels could act as powerful inputs to the standard classifiers. Later on, researchers tested and verified this idea on actual hardware by developing an end-to-end workflow. Havlicek et al. [17] built a quantum kernel estimator that feeds a variational quantum classifier and a classical SVM on a superconducting processor. A recent research by Huang et al. [18] showed through demonstration on today's quantum hardware that a quantum approach can significantly reduce the number of physical experiments needed for particular "learning from experiments" tasks. Figure 2 shows the workflow of a hybrid QML framework. The workflow begins with preprocessing and encoding of the classical data for execution on a quantum circuit. The output of the quantum circuit is measured and post-processed classically and the output is fed into a classical optimization loop that updates the quantum circuit's parameters iteratively to train the model.

Broadly, the QML models can be divided into four categories - supervised quantum kernel methods, variational or discriminative models, quantum generative models, and quantum reinforcement learning models. Quantum-kernel

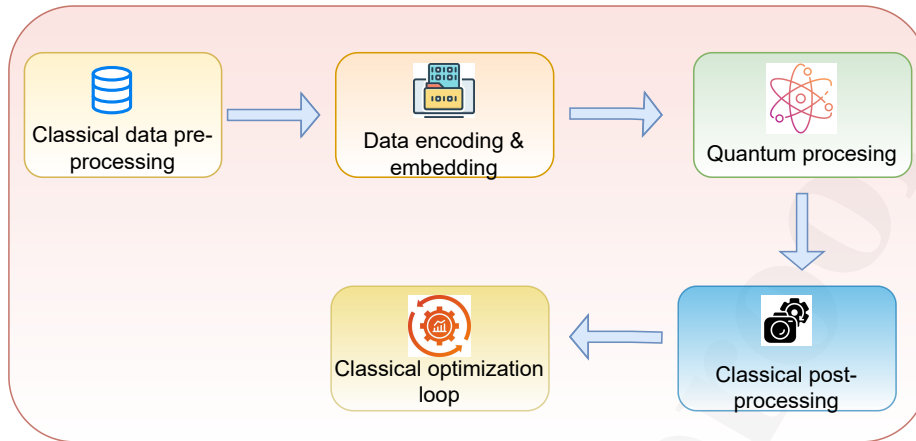


Figure 2: Workflow of a hybrid quantum-classical machine learning framework

methods embed classical data x into quantum states $|\phi(x)\rangle$ and rely on the similarities that are estimated as inner products (kernels) by the model. The quantum kernel methods make use of a quantum processor only to compute the kernel matrix $K_{ij} = |\langle \phi(x_i) | \phi(x_j) \rangle|^2$ and a classical model is used to handle the training. Variants of quantum-kernel methods differ in the feature map (amplitude, angle, or problem-structured encodings) and in terms of how the kernels are estimated. With these methods, the training remains a convex classical optimization, and hence they are attractive when the data are scarce and the decision boundaries are subtle. They also provide better audit trails (margins and support vectors). Variational or discriminative models (parametrized quantum circuits, PQC) have their quantum part as a trainable circuit $U(\theta)$ whose outputs (bitstrings, expectation values) are fed to a loss that is minimized by a classical optimizer. Cerezo et al. [19] provided a comprehensive review on variational quantum circuits, explaining the standard loop of prepare-measure-update, known training pathologies, design choices for ansätze, and mitigation strategies for near-term devices. Geometry-aware optimization and data re-uploading mechanisms increase the practical expressivity of PQC without the need for deep circuits. PQC are preferred in practicality when full end-to-end quantum training of the model is desired (instead of a “kernel-only” role). Quantum generative models help to learn data distributions using quantum processes [20]. Quantum Generative Adversarial Networks (QGANs) train a quantum generator against a quantum or classical discriminator. Quantum Boltzmann machines (QBM) model thermal-like distributions. Lloyd et al. [21] present an information-theoretic perspective on quantum generative modeling. Generally, quantum generative models are utilized for tasks where sampling structure (interference effects) is central. They are also used in places where classical scoring can be integrated with the evaluations of sample quality or model likelihoods.

Quantum reinforcement learning (QRL) uses quantum components inside policy/value approximators or uses

quantum subroutines within the normal loop of reinforcement learning. Jerbi et al. [22] present a survey covering algorithm templates (PQC policies, exploration primitives, value estimation), small-scale demonstrations, and open issues. The survey focuses on where quantum elements plausibly fit and where the QRL models might be better than the variational classifiers or kernel methods. Across all the above categories, the practical QML models remain hybrid - quantum hardware is used for well-defined subroutines like kernel estimation and circuit evaluations, while data handling and optimization remain classical.

QML models have multi-fold advantages compared to their classical counterparts. As discussed before, quantum kernel mapping makes the feature space more expressive, which is hard to achieve with classical methods. PQC follow analytic gradient rules, which enable the estimation of exact derivatives on real devices without the need to resort to fragile finite differences. Schuld et al. [23] showed that this makes gradient-based training practical on hardware and helps stabilize optimization pipelines. Some quantum architectures encode symmetry and locality directly into the model, benefiting from built-in inductive bias from symmetry and physics. For example, Quantum Convolutional Neural Networks (QCNNs) use shallow and hierarchical circuits that reflect renormalization ideas [24]. QCNNs are proven to classify cases of matter with logarithmic complexity and good noise tolerance. Apart from the raw accuracy, the theory for generalization that represents how well a trained model will perform on new data is an emerging research direction in QML. They utilize quantification techniques like Fisher information/effective dimension, giving insights into when a quantum model is likely to generalize well or overfit [25].

Being a relatively newer field, QML is not without significant challenges that need to be overcome to realize its full range of benefits. The practical guidance in the current variational circuits literature is to prefer problem-structured and shallow circuits with careful initialization and optimization. This is due to the barren

plateau problem, where large and unstructured PQCs have gradients that vanish as systems get larger. This makes training ineffective [19]. Noise and resource overheads are another notable challenge in current-day QML models. Error mitigation can reduce bias for short circuits, but at the cost of extra sampling [26]. Recent research showed that for a particular set of quantum linear-algebra-based learners, quantum-inspired classical models achieve similar performance to that of quantum models under realistic input models. This denotes the experimentation in the lines of dequantization. Further, the comparisons with classical baselines can be skewed due to weak classical baselines and inconsistent processing. Standard datasets/splits, full reporting of noise and hardware, and matched feature budgets could be some of the solutions for the above-mentioned issue.

It is also important to acknowledge that QML has not consistently outperformed classical methods, as of the current state of QML research. The dequantization results by Tang et al. [27] showed that quantum-inspired classical algorithms can match the runtime advantages that were previously attributed to linear-algebra and quantum recommendation-based learners. Huang et al. [18] demonstrated that the classical ML models on the same data achieve comparable accuracy to QML models, with the quantum advantages being limited to specific settings where data comes from quantum processes. It has also been noted that the noise-induced performance degradation has resulted in quantum classifiers underperforming optimized classical baselines on several benchmarks on current NISQ hardware [28]. These findings indicate that the claimed QML advantages are task-dependent and conditional on hardware maturity.

IV. CORE MECHANISMS AND ALGORITHMIC PRINCIPLES

A. Quantum-enhanced Learning Pipelines

Quantum-enhanced learning pipelines fit into today's control decision systems, where they can be used only when they are required for certain subtasks. They can integrate seamlessly into the current MC control and decision systems, while leveraging quantum processors only where they provide significant representational or computational advantages. The mission critical domains share common requirements - extracting meaningful insights and patterns from small & high-dimensional datasets, maintaining predictable and stable model training under hardware constraints, and achieving required expressivity with shallow circuits. In this section, we try to map these requirements to quantum machine learning methods.

In applications such as satellite telemetry anomaly detection and propulsion systems fault diagnosis, the classical ML models face limitations in handling limited samples of high-dimensional data. Quantum feature mapping projects the classical data into a higher-dimensional Hilbert space, which enables the separation of complex patterns even with fewer samples. A kernel, denoted by $k(x, x')$, serves as a measure of similarity between quantum states, when a classical input x is encoded into a quantum state.

$$k(x, x') = |\langle \phi(x) | \phi(x') \rangle|^2$$

In MC contexts, this kernel helps in the detection of subtle separations in high-dimensional sensor or telemetry data, thus distinguishing normals and anomalous satellite readings, even in the case when labeled fault samples are scarce.

Schuld et al. [16] proposed that a simple encoding of the inputs into quantum states works like a non-linear feature map. The authors provide explicit constructions for such kernels. Havlíček et al. [17] implemented two quantum-classical hybrid pipelines on a superconducting processor, using a quantum variational classifier and a quantum kernel estimator combined with classical SVM. The results of this study confirm that it is possible to realize quantum-enhanced feature spaces on existing hardware to support data-scarce MC applications like integrity monitoring in safety-critical networks. Stability and predictability of training dynamics are another critical aspect in MC systems. The classical stochastic optimizers may exhibit slow convergence and instability in bad loss landscapes, which is unacceptable in real-time and embedded environments. Quantum Natural Gradient Descent (QNGD) mitigates this issue by incorporating the quantum state space geometry into parameter updates. Stokes et al. [29] demonstrated that QNGD allows the optimizer to move in the direction of the steepest descent with respect to the quantum information geometry. In MC applications, QNGD provides a theoretically grounded approach for stable training in long-duration missions and adaptive control in dynamic environments. In some of the resource-constrained MC environments, like edge devices on autonomous platforms, efficient circuit designs that meet decoherency and latency limits, and at the same time don't sacrifice model capacity. Pérez-Salinas et al. [30] proposed a data re-uploading technique that satisfies the above condition. This technique embeds the same input several times within a shallow quantum circuit, which enables expressive decision boundaries with minimal depth. Data re-uploading techniques provide a low-latency solution for multiple inference tasks on satellites, autonomous vehicles, and UAVs.

Since the current quantum hardware lacks fault tolerance, researchers have proposed hybrid quantum-classical frameworks, which are a more practical near-term option for MC deployment. Variational Quantum Algorithms (VQAs) follow this design philosophy that delegates the initial quantum state preparation and the final measurement to the quantum device and keeps optimization and loss computation on a classical processor. Cerezo et al. [19] provided detailed architectures, mitigation approaches, and convergence analysis of VQAs, which are compatible with noisy-intermediate scale quantum hardware. This hybrid structure of VQAs allows safety interlocks and time-critical I/O to be under classical control while asynchronously invoking quantum routines. This configuration is ideal for auditable operations in energy optimization and disaster management.

One of the notable requirements for MC integration is the credible experimental validation of learning performance on physical quantum devices. Huang et al. [18] presented theoretical and experimental demonstrations of "quantum advantage" in learning tasks. The authors validated on 40 superconducting qubits that executed around 1300 quantum gates. They reported that the quantum systems could learn

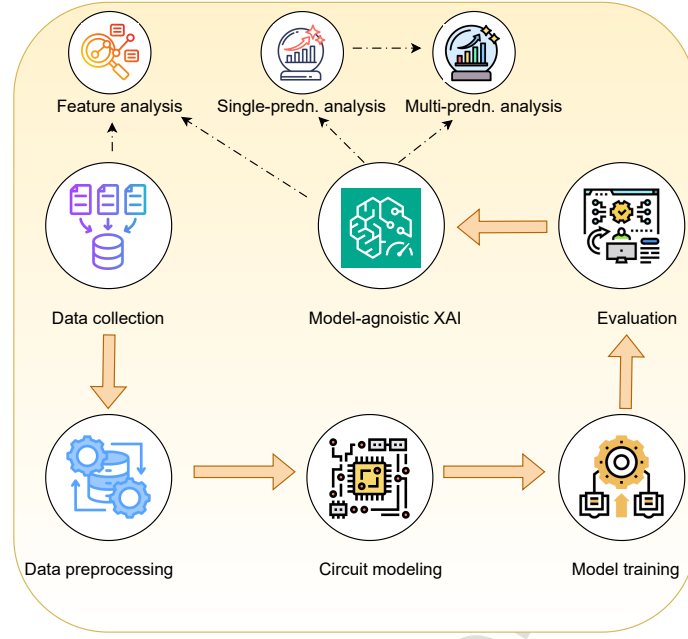


Figure 3: Workflow of a model-agnostic explainable quantum machine learning framework

from an exponentially smaller number of experiments. These findings strengthen the fact that meaningful workloads can already be executed on real hardware, which supports calibration, diagnostics, and adaptive control in space missions and aerospace.

B. Quantum Uncertainty, Robustness, and Fault Tolerance

Predictable behavior under uncertainty is a critical need of MC systems such as defense decision loops, spacecraft control, and autonomous air-traffic coordination. Asymptotics and heuristics are generally employed in classical systems to estimate uncertainty. These can lead to miscalibration in distribution-shifted or small-data regimes. In quantum computing, a technique called Quantum Conformal Prediction (QCP) uses conformal prediction along with the intrinsic randomness of quantum measurements and hardware noise to produce valid error limits. Park et al. [31] noted that due to device and measurement noise, their quantum models yield multiple random decisions per input. The authors leveraged randomness to define prediction sets that capture the uncertainty of the model provably, providing finite-sample coverage guarantees for both regression and classification. The authors also performed experiments on existing quantum hardware, along with simulators, to confirm the calibration guarantees. The basic building block for uncertainty is the variance of an observable parameter A , measured on a state ρ :

$$(\Delta A)^2 = \langle A^2 \rangle - \langle A \rangle^2, \quad \langle A \rangle = \text{Tr}(\rho A)$$

where $\text{Tr}(\cdot)$ denotes the trace.

Another requirement in MC systems is the robustness to device noise during inference and training. As of current research, full-scale quantum error correction (QEC) is an

asymptotic solution. Fowler et al. [7] provided quantitative fault-tolerance estimates for surface-code logical qubits, decoding, and operations, indicating their practicality for large-scale systems. Error mitigation naturally occurs with learning/variational loops. Two popularly analyzed techniques for mitigation are probabilistic error correction and Zero Noise Extrapolation (ZNE). ZNE evaluates quantum circuit at amplified noise levels and extrapolates to zero noise [26]. ZNE fit can be represented as

$$E(0) \approx \sum_{i=1}^m c_i E(\lambda_i), \quad \text{s.t.} \quad \sum_i c_i = 1$$

Where $E(\lambda_i)$ is the measured expectation at λ_i (amplified noise factor) and c_i represent extrapolation coefficients that are chosen to cancel low-order noise terms. Li et al. [32] integrated the error mitigation with Variational Quantum Classifier (VQC) and demonstrated improvement in the learning.

Resilience to adversarial effects and perturbations, combined with the tools for formal robustness verification, is another requirement in MC systems. Randomized encodings and quantum noise can certify robustness bounds of quantum classifiers (against adversarial examples) on the defensive side. Huang et al. [33] derived robustness bounds and connected them with the notion of differential privacy. Lin et al. [34] proposed VeriQR, which provides an automated formal verification workflow with adversarial example discovery, local/global robustness checks, and adversarial training. VeriQR can enable pre-mission certification (verification under disturbances) and in-mission monitoring that triggers fallbacks at violation of verified margins.

C. Quantum Explainability and Verification

MC deployments in disaster-response logistics, satellite operations, safety-critical transport, etc., require that the output of quantum AI be not only accurate, but also interpretable, trustworthy, and verifiable against perturbations and implementation errors.

Formal robustness verification oriented to quantum classifiers can potentially benefit the trustworthiness of QAI decisions in MC contexts. VeriQR framework [34] (also discussed in Section IV-B) provides complete and sound algorithms for robustness verification (local and global) of QAI models under realistic device noise. VeriQR can automatically find adversarial examples and perform adversarial training to improve robustness. These properties provide concrete verification proofs that are comparable to classical safety cases. Caro et al. [35] proposed protocols for classical verification of quantum learning based on interactive-proof frameworks. These protocols allow the classical clients to verify an untrusted quantum server executing a delegated learning task in terms of correctness of execution. This establishes correctness guarantees for remote QAI. Hence, the works by Lin et al. [34] and Caro et al. [35] give evidence of the feasibility of delegated-execution verification and robustness checks with current QAI formulations, thus addressing the audit requirement of MC deployments.

Recently, several research papers have been adapting feature-attribution ideas of classical ML explainability, like LIME (Local Interpretable Model-agnostic Explanations) and SHAP (Shapley Additive Explanations), to QAI models, including quantum kernels, parametrized quantum circuits, and QSVMs. Figure 3 presents the workflow of a model-agnostic explainable quantum machine learning framework. It shows how data is collected and processed to build and train a quantum circuit, which is evaluated using model-agnostic explainable AI methods. The model-agnostic explainable evaluates the trained model and feeds its interpretations to refine feature analysis and analyze both single and multi-prediction outcomes. In this framework, Kadian et al. [36] proposed ExQUAL (Explainable Quantum Classifier), integrating LIME and SHAP into Pegasus QSVM to produce both global and local explanations for binary and multi-class tasks. This shows that the classical Explainable AI(XAI) methods can be combined with trained quantum models and quantum feature maps to obtain human-interpretable attributions. Heese et al. [37] apply the concept of Shapley values to quantify the importance of groups of gates in a QAI model for attaining specific goals. The authors proposed that the resulting attributions can be seen as explanations for why a particular circuit works well for a given task, which further improves the understanding of how to construct both parametrized and variational quantum circuits. The authors demonstrated the benefits of the proposed framework for transpilation, classification, optimization, and generative modeling with experimental evaluation on both simulators and superconducting hardware devices. Kottahachchi et al. [38] proposed QRLaXAI (Quantum Representation Learning and Explainable AI), combining a quantum autoencoder

with a variational quantum classifier and incorporating empirical and theoretical explainability for image data/ The authors leveraged visual explanations of LIME and analytical insights of SHAP models to attain a deeper understanding of the model's decision-making process based on output. It is important to acknowledge here that, while these adaptations of classical XAI techniques to QAI models are a necessary step, they are not sufficient for MC certification. It remains an open question whether LIME- and SHAP-derived explanations of quantum models constitute acceptable evidence under standards like IEC 61508 and DO-178C, as no certification authority has yet evaluated these quantum-specific explainability components. Further, there is also scope for designing explainability metrics specific to quantum machine learning models [39].

V. APPLICATIONS

Quantum AI has multifarious applications in mission critical systems – some of them are highlighted in Figure 4 along with their discussion below.

A. Aerospace and Space Systems

Aerospace planning tasks like trajectory design, de-confliction, and constellation scheduling are treated as large combinatorial problems. Several research works show that these tasks can be mapped to quadratic unconstrained binary optimization (QUBO) and solved by exploiting quantum heuristics, such as variational approaches and annealing. QUBO is a combinatorial optimization formulation where binary decision variables are optimized over a quadratic objective without any explicit constraints. In QUBO, all problem constraints are absorbed into the objective as penalty terms. QUBO is natively compatible with variational quantum algorithms such as QAOA and quantum annealing hardware. Stollenwerk et al. [40] constructed a conflict-graph formulation for strategic air traffic de-confliction. The authors showed how wind-optimal multiple flights can be mapped to QUBO and further evaluated the problem instances on classical baselines and D-Wave hardware, thus demonstrating an end-to-end encoding and solution workflow. For spacecraft trajectory optimization, Grossi et al. [41] transcribed continuous-time dynamics into a binary QUBO and proposed to solve interplanetary transfers using quantum annealing. The authors also dealt with discretization and constraint-handling steps that are required to obtain feasible trajectories. Makhanov et al. [42] proposed a modular hybrid model for flight-path optimization and analyzed resource constraints in quantum backends and classical simulators, clearly indicating where the quantum subroutines can be fitted in present-day pipelines. In all these formulations, the standard QUBO takes the form:

$$E(\mathbf{s}) = \mathbf{s}^T Q \mathbf{s}$$

where $\mathbf{s} \in \{0, 1\}^n$ are decision variables (such as discrete heading, timing, assignment choices) and $Q \in \mathbb{R}^{n \times n}$ encode penalties and costs derived from hard constraints and the



Figure 4: Applications of Quantum AI in Mission critical systems

objective of the mission, such as slew limits, fuel and separation minima.

Aerospace operations demand quantified uncertainty and reliable fault detection with telemetry systems. Quantum anomaly detection methods provide procedures to detect variations from the normal behavior after encoding the classical data into quantum feature spaces. Liang et al. [43] proposed a quantum anomaly detection method based on density estimation and multivariate Gaussian distribution. The authors constructed both algorithms using the standard gate-based model of quantum computing. They also presented a quantum procedure for efficiently estimating the determinant of any Hermitian operator.

Multi-agent decision-making under partial information is central in autonomous spacecraft control and coordinated swarms. Chen et al. [44] introduced a quantum multi-agent reinforcement learning framework (QMARL), which functions based on quantized state/action representations and reports coordination and navigation experiments in multi-robot settings. The authors provided implementation details of the training loop for decentralized agents and PQC-based policy modules. Although the aerospace swarm-flight optimization is an engineering task, the study of Chen et al. provides a concrete algorithmic template (state encoding, policy circuit design, and training protocol) for migrating multi-agent navigation/avoidance to quantum-enhanced controllers. Further, the classical literature [45] on distributed cooperative control for spacecraft swarms provides the connectivity models and the constraints for quantum reinforcement learning (QRL) models to follow for collision-avoidance and connectivity maintenance. This literature also clarifies

the interface between the learning-based policies and the guidance laws.

B. Defense and Autonomous Systems

Defense autonomy combines multi-sensor perception, mission planning, and adaptive control under strict timing and safety constraints. The planning layer in these autonomous systems is generally reduced to NP-hard scheduling and routing, and the autonomy part balances fuel/time budgets, online re-planning under uncertainty, and collision avoidance [46].

Under quantum optimization for mission routing and tasking, Mori et al. [47] formulated a real routing problem as a QUBO and solved it using quantum clustering and annealing-based pre-processing. The paper presents an end-to-end encoding and reports the solution quality on practical instances, thus illustrating how logistics-type planning in defense autonomy can be mapped onto annealing hardware. The authors demonstrate that the annealing-based solvers can easily handle real constraints like capacities and time windows within hybrid pipelines. For sensor-driven navigation and collision avoidance, Sinha et al. [48] developed quantum deep reinforcement learning agents that included a policy based on PQCs. The authors evaluated the proposed framework on autonomous navigation tasks and reported improved training stability (less sensitivity to initialization of the weights), a higher average cumulative reward compared to classical baselines, and explicit deterioration under injected quantum noise, thus quantifying the robustness tradeoffs. In

the proposed framework, the optimized target is the discounted return:

$$J(\theta) = \mathbb{E} \left[\sum_{t=0}^T \gamma^t r_t \mid \pi_\theta \right]$$

where r_t is the reward at time t , γ is the discount factor, π_θ is the PQC-parametrized policy with parameters θ , and T is the episode horizon. In defense autonomy, a requirement is that the policies must retain performance under real-world disturbances and hardware noise; hence, the work of Sinha et al. is directly relevant because of its noisy-simulation study.

For model reliability under device noise, Heyraud et al. [49] characterized how dephasing and dissipation impact quantum-kernel classifiers and support vector machines. The authors quantified performance loss as a function of circuit depth and noise parameters, and discussed regularization and mitigation strategies. The results of this work provide guidance on operating regimes (where kernel-based QAI remains reliable) for defense sensor-fusion stacks adopting quantum kernels.

The upcoming Agentic AI (AAI) technology can potentially benefit the multi-UAV defense applications in several ways. Khowaja et al. [50] demonstrated that AAI systems can be used for autonomously analyzing multimodal sensor data, dynamically allocating resources, and coordinating responses in real time, thus reducing initial response times by over 60% compared to rule-based systems. The integration of QAI techniques can potentially accelerate the capabilities of AAI-based UAV systems by enabling exponentially faster optimization of resource allocation and threat prediction across vast and complex operational environments that might overwhelm classical computing approaches.

C. Cybersecurity and Critical Infrastructure Protection

Critical infrastructure networks like power, water, and transportation require timely detection of advanced threats and cryptographic resilience in an increasing manner. Quantum Key Distribution (QKD) can be used in utility fiber networks for high-assurance keying; further, quantum keys can be used for authentication of grid communications. QKD is a cryptographic protocol that exploits quantum mechanical properties like the no-cloning theorem and measurement disturbance in order to generate a shared secret key. The security of this key is guaranteed by the laws of physics rather than computational hardness assumptions. Intrusion Detection Systems (IDS) in cybersecurity are increasingly relying on machine learning; the quantum machine learning tools can potentially replace them in areas or situations where a significant quantum advantage might be reaped [51]. Secure collaboration across organizations demands privacy-preserving learning, and the recent quantum homomorphic encryption (QHE) based decentralized learning frameworks provide data privacy guarantees that are computation-theoretic.

Alshowkan et al. [52] made use of QKD protocols in the authentication of smart grid communications, which are executed on a deployed electric-utility fiber network. The authors created a software stack that ingests keys from a QKD system and uses them to generate message-authentication

codes for control links in machine-to-machine scenarios. Thus, the proposed framework strengthens operational traffic with keys that have physical-layer security properties. Kumar et al. [53] proposed QuIDS, an intrusion detection system based on a quantum support vector machine. The authors proposed to detect IoT attack classes using small training sets of low-dimensional features. The authors demonstrate competitive results against benchmarks in IoT intrusion detection systems. The decision function of QSVM follows the standard large-margin form:

$$f(x) = \text{sign} \left(\sum_{i=1}^m \alpha_i y_i k(x_i, x) + b \right)$$

where x is the feature vector, (x_i, y_i) for $i = 1$ to m are the support vectors x_i with labels y_i , $k(\cdot, \cdot)$ is the quantum kernel, b is the bias, and α_i are the learned dual weights. In MC cybersecurity settings, this decision function helps in classifying the incoming network flows as benign or malicious. The quantum kernel can improve the separation of attack signatures from classical feature spaces.

Kalinin et al. [54] also evaluated QSVMs and quantum convolutional neural networks (QCNNs) against classical IDS systems on standard traffic datasets. The authors encode the flow/packet features into PQCs and compare performance under matched training budgets. The paper not only shows the construction of quantum feature maps for network security data, but it also quantifies cases like limited data, where quantum kernels are competitive with classical detections.

Li et al. [55] proposed a QHE based federated learning framework for cross-operator collaboration without the need to expose raw data. In this scheme, the clients encrypt quantum data and delegate tracing to a server, and get outputs that have computation-theoretic privacy [56]. The authors demonstrated that the proposed framework reduces communication complexity significantly compared to blind quantum computing approaches.

D. Energy and Smart Grids

The operation of smart grids require reliable and fast decisions for stability and situational awareness. The classical frameworks are stressed mainly due to the fact that the core optimizations are computationally hard (AC optimal power flow is strongly NP-hard), as shown in research [57]. This complicates real-time convergence and guarantees of classical solvers. The variation in solar and wind energies provides steep load ramps that require multi-interval flexibility from conventional resources [58]. Recent research explored quantum reinforcement learning to assist stability control and contingency analysis in smart grids. Peter et al. [59] applied QRL for power grid security assessment, particularly for combinatorially difficult contingency analysis problems (where classical RL fails in terms of computational cost). The proposed QRL framework helped scale by improving computational efficiency and agent proficiency. The authors demonstrated a proof-of-concept by using QC for RL agent simulation and training. The proposed hybrid agent runs on quantum hardware based on IBM's Qiskit runtime. The

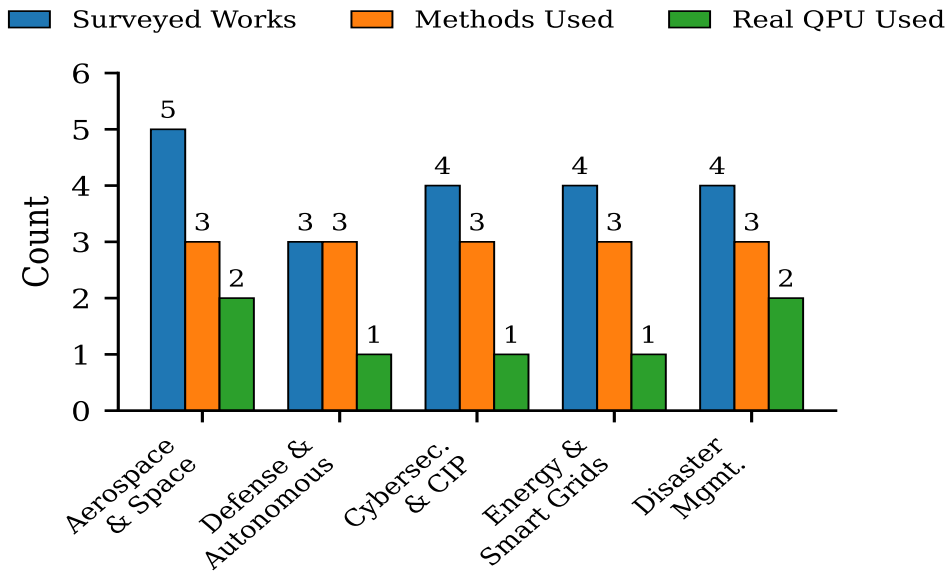


Figure 5: Comparative Survey Summary Across MC Application Domains

optimized target in the proposed framework is the discounted return.

Unit commitment (UC) and Optimal Power Flow (OPF) remain critical for scheduling and resource allocation in smart grids. Unit commitment is a short-term scheduling optimization problem that decides which of the generating units should be turned off or on over a planning horizon, which typically ranges from 24-168 hours, with the aim of meeting the forecast demand at minimum cost. It turns out to be a large-scale mixed-integer optimization problem after incorporating operational constraints like minimum up/down times, ramp rates, and startup/shutdown costs. Optimal Power Flow is a steady-state network optimization problem that decides the optimal operating point of a power system with minimization of an objective (generally, generation cost or transmission losses), while satisfying both physical laws (Kirchhoff's laws) and operation constraints (line thermal limits, voltage limits).

Salgado et al. [60] developed a hybrid quantum-classical algorithm to approximately (efficiently) solve the UC problem in polynomial time (the UC problem is NP-hard). The authors decomposed the UC problem into two subproblems - a QUBO and a quadratic optimization problem. The authors made use of QAOA to determine the optimal unit combination, along with classical methods to identify individual unit powers. The proposed framework reduced the number of iterations required for QAOA convergence and also improved the solution accuracy.

For fault prediction and anomaly detection, Ajagekar et al. [61] proposed a QC-based deep learning framework. The deep learning stack consists of a conditional Boltzmann machine for feature extraction and deep neural networks for classification. The authors addressed the computational challenges of training

deep learning models with QC-based training methodologies. The authors reported improved computational efficiency and faster response time with the proposed hybrid model on their specific benchmark, though these gains remain to be validated at production scale.

E. Disaster Management

Disaster management requires logistic planning, early warning, infrastructure restoration, and rapid evacuation under severe uncertainty. The classical formulations with evacuation scheduling and routing, emergency vehicle routing, and post-disaster restoration are treated as NP-hard, which leads to heuristic screening or approximations in time-critical settings [62]. In this landscape, QAI can potentially provide significant advantages with quantum-accelerated restoration scheduling, quantum optimization for logistics/evacuation, and quantum-enhanced early warning systems.

Liu et al. [63] conducted a performance evaluation of tsunami evacuation route planning by solving a combinatorial optimization problem on quantum and non-quantum annealers. The authors modeled the evacuation problem mathematically and converted it into QUBO form. The authors evaluated the performance of four annealing machines, including D-Wave Neal, D-Wave Advantage, Fixstars Amplify Annealing Engine, and Vector Annealing. They conducted an evaluation from three aspects: the solution quality, time to solution (TTS), and the maximum number of evacuees. Fu et al. [64] proposed a hybrid quantum-classical approach for coordinated post-disaster restoration (CPR). The authors established the CPR model with the consideration of power distribution system (PDS) operation, topology configuration, microgrid operation, and repair crew dispatch. They made use of binary expression and integer slack methods to transform the

Companies involved	Topic of collaboration	Brief description on the work	Stage
SandboxAQ & United States Air Force (USAF)	AI-assisted magnetic navigation (Quantum)	Testing AI-assisted navigation using quantum magnetometers as a primary guidance method (without GPS)	Field demo
Air Force Research Laboratory (AFRL) & QC Ware	Quantum Machine Learning (Q-means)	Exploring quantum clustering (q-means) to group unmanned aircraft flight paths and infer intent.	PoC
Airbus, QC Ware & IonQ	Aerospace optimization (Hybrid Quantum)	Running hybrid quantum programs for cargo loading, routing, and other optimization problems.	PoC
IBM & Raytheon	Quantum computing and AI for defense	Collaborating to co-develop solutions that integrate QC and AI particularly for applications in the defense sector.	Exploration
Terra Quantum & (Utility partners)	Energy time-series forecasting (Hybrid Quantum)	Developing hybrid quantum models to improve short-term predictions for energy metrics for grid balancing.	PoC
D-Wave, TNO & Quantum Quants	Grid optimization (Hybrid Quantum)	Developing a hybrid quantum constrained quadratic model solver for energy grid redistribution and partitioning.	PoC
IBM & E.ON	Utility operations and forecasting (Quantum/QML)	Exploring quantum methods for complex utility operations.	Exploration
IonQ & Hyundai	1. Battery chemistry (Quantum) 2. Machine perception (Quantum ML)	1. Using variational quantum eigensolver to study lithium compounds. 2. Experimenting with QML for machine perception.	PoC

Table I: A summary of industrial practice of QAI in MC systems (PoC: proof-of-concept without production deployment; Field demo: tested in an operational environment; Exploration: early-stage collaboration without published results)

constrained problem into QUBO. D-Wave's direct quantum processing unit solvers and hybrid solvers of Leap quantum cloud service are employed to perform a comparative analysis of different calculation methods. Sankaradass et al. [65] proposed QSVM and QNN-based models for early warning systems for environmental disaster prediction. The authors tested their proposed models on two benchmark datasets, and the findings reveal that the proposed QAI models are accurate than the classical ML models. Kobayashi et al. [66] outlined an evacuation planning stack that combined the D-Wave quantum annealing engine with the NEC SX-Aurora vector supercomputer. They presented a high-level workflow of evacuation models, which consists of the initial QUBO mapping and partitioning of computation between the annealer for the last-stage optimization and classical vector hardware. They also provide a concrete systems blueprint for integrating HPC evacuation simulators with annealing-based optimizers.

Figure 5 presents a bar chart quantifying the surveyed works, methods used, and real QPU used across five MC domains. The surveyed works on aerospace used three methods - QUBO/annealing, anomaly detection, QMARL; works on defense used QUBO/annealing, QRL, quantum kernels; works on energy used QRL, QAOA, quantum boltzmann machines; and works on disaster management used

QUBO/annealing, QSVM, and QNN.

VI. INDUSTRIAL PRACTICE

Several companies are actively exploring and developing proofs-of-concept (PoCs) on application of QAI methods to solve challenges in MC systems, particularly in aerospace and defense. They are using simulated quantum computers and current-day NISQ quantum computers. A summary of industrial efforts in this area is noted in Table I.

SandboxAQ [67], in partnership with the United States Air Force (USAF), is testing AI-assisted magnetic navigation that uses quantum magnetometers to navigate, without the need for GPS. The USAF flew the system from Joint Base Charleston, where magnetic navigation was used as the primary guidance method during the flight, thus reporting a real-time demo. Air Force Research Laboratory (AFRL) partnered with QC Ware to explore quantum machine learning clustering, termed q-means, to group unmanned aircraft flight paths and infer intent. This is an early research towards situational awareness and autonomy. Airbus, in partnership with QC Ware and IonQ, runs quantum programs on cargo loading, routing, and other related aerospace optimization problems [68]. The aim of this project is to examine how hybrid quantum approaches could surpass or complement

existing High Performance Computing (HPC) heuristics. IBM and Raytheon (an Air Force and defense tech company) forged a collaboration to co-develop solutions that combine quantum computing and AI for the defense sector [69]. Terra Quantum is developing hybrid quantum models for energy time series forecasting, such as steam flow from power plants and solar output [70]. The aim of the project is to improve the short-term predictions for grid balancing and operations planning with utility partners. D-Wave is partnering with TNO and Quantum Quants to develop hybrid quantum constrained optimization on redistribution and grid partitioning scenarios. Their hybrid constrained quadratic model solver reported to find high-quality solutions within fixed time budgets, compared to several classical baselines [71]. IBM is collaborating with E.ON (Europe’s largest energy company) to explore quantum methods for complex utility operations and forecasting at scale, plausibly using QAI techniques. The main focus is planning and optimization under renewable energy variability, which is being dealt with as exploration rather than a production rollout. IonQ and Hyundai are working on two separate research threads relevant to future transportation systems [72]. Firstly, they are planning to develop next-generation battery chemistry technology by using Variational Quantum Eigensolver (VQE) to study lithium compounds, which also aims at improving the safety and performance. The companies are also pursuing quantum ML experiments for machine perception, starting with road-sign classification, and moving towards 3D point-cloud object detection.

VII. A MODEL FOR QUANTUM RESOURCE MANAGEMENT AND SCHEDULING SYSTEM

There has been extensive research in the classical/high-performance computing cloud scheduling and resource management, but quantum cloud computing, being nascent, lacks in this. It is not possible to completely use the classical scheduling and resource management systems in quantum cloud computing, due to the different nature of the quantum tasks. Hence, we propose a framework to handle the unique properties and constraints of quantum computing in resource management and scheduling systems.

We assume the following deployment conditions for the framework. (i) a heterogeneous pool of NISQ-scale QPUs (up to a few hundred qubits) are accessible via cloud APIs, (ii) data regarding real-time calibration like gate fidelities and coherence times is available from each QPU, (iii) the quantum jobs cannot be checkpointed or preempted in between execution, and (iv) classical pre-processing, post-processing, and orchestration are co-located with the scheduling service.

Figure 6 shows the proposed software framework. It begins with API Gateway and Authorization component, which acts as a secure front door for all the interactions of users. Apart from the standard authentication and quota management, this component is particularly proposed to ingest complex quantum task definitions, including serialized objects from frameworks such as Cirq or Qiskit and the circuits in QASM format. This component also accepts other user-defined

preferences and constraints, including the minimum acceptable 2-qubit gate fidelity, required qubit count, the maximum acceptable queue wait time, and the priority of the task. All of these metadata are essential for downstream decision making.

After validation, a job will be passed to a `Job Queue`, which is a scalable and persistent buffer that separates the intensive scheduling process and high-volume submission of tasks. Although this job queue is functionally similar to classical queues, it will be optimized to store and prioritize jobs based on their quantum metadata. The `Preprocessing & Analysis` service retrieves jobs from the job queue to execute a “quantum-aware” inspection. This component estimates several fundamental properties of the user’s circuit by performing static analysis, including the circuit depth, the exact qubit count, and the number of high-error-rate 2-qubit gates. An important function of this component is to generate an accurate estimate of the quantum execution time and compare it against known coherence times, which helps in flagging jobs that are too deep to be successful. Based on this analysis, it assigns tags to jobs (like type-deep circuit, require-high fidelity) that are very helpful to the scheduler.

Like in classical scheduling, the system requires a real-time view of the hardware, which needs to be supported by the `QPU Hardware Monitor`. Compared to a classical CPU monitor, this component is very complex and it continuously polls all the QPUs to track both the static properties like qubit count and topology, and the other dynamic calibration-dependent properties. This monitor also produces data such as 1-qubit and 2-qubit gate fidelities and current gate execution times. The monitor also reports the QPU’s state (online, offline, or recalibrating), making sure that the scheduler never uses inaccurate or stale data. The `Scheduler` component is the most important component that takes in the data from the hardware monitor and decides the jobs to run on different QPUs. It also takes into special consideration quantum jobs. For example, the quantum jobs cannot be paused and resumed later with the help of some kind of checkpointing mechanism. Firstly, a feasibility filter narrows the QPU choices by finding all those that are available and that meet the job’s requirements. Next, a policy engine will be used to rank the feasible QPUs to select the best. In this work, we adapt a classical HPC strategy - shortest job first, which uses the estimated execution time to maximize throughput. Other classical policies like round robin, shortest remaining time first, and multilevel feedback queues could also be adopted based on the requirement. The scheduler also considers quantum-aware policies like best-fit fidelity that try to match a job to a QPU just above the required fidelity. This helps in saving the high-performing and rare QPUs for the jobs that require them the most. For n queued jobs and k available QPUs, the feasibility filter evaluates each job against each QPU, resulting in $O(nk)$ complexity. The subsequent policy ranking sorts feasible candidates in $O(n \log n)$ for priority-queue-based policies like shortest job first.

The jobs will be sent to the `Transpiler` service after the scheduler makes its selection, which performs heavy classical computation to translate the user’s logical circuit into a physical circuit for the specific QPUs. Since there are

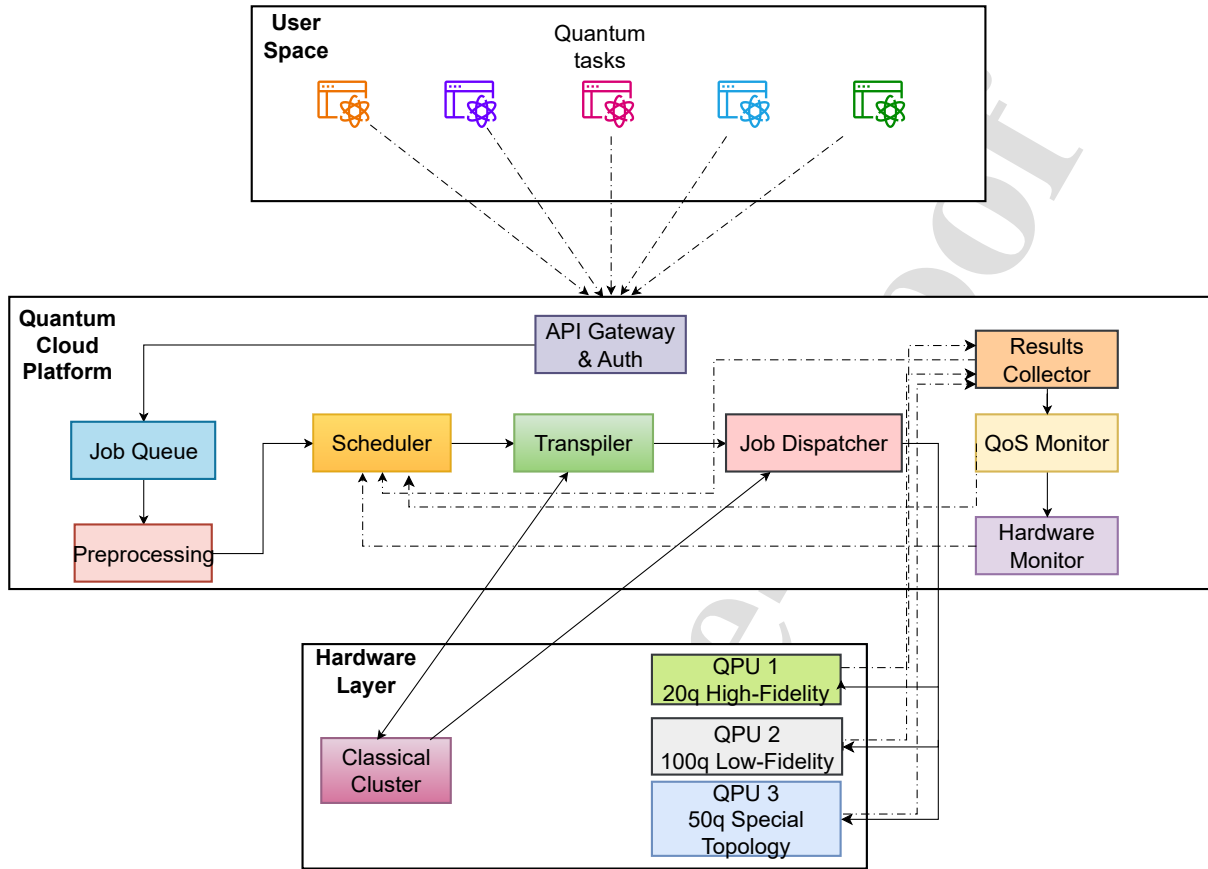


Figure 6: Proposed Quantum Cloud Scheduling and Resource Management Framework

different types of quantum hardware, the transpiler should be able to translate the user's code to the required QPU code. It is possible that the conversion process adds error and depth, so a sophisticated scheduler may use a trial transpilation to select the QPU that gives the lowest swap overhead. After this, the final, machine-ready program will be handed over to the Job Dispatcher, which opens a connection to a QPU and executes the program for a particular number of shots. The raw results will be collected by the Results Collector & Post-processing service, whose important quantum-specific role is to apply error mitigation techniques, analogous to the classical post-processing steps that use calibration data to correct for known measurement errors, before handing the final answer to the user. A QoS Monitor and Feedback loop observes the entire system's performance, tracking metrics like the job wait time that help the processing service improve its execution time estimates for future jobs, and also the hardware monitor by indicating any QPU that might be performing poorly.

Table II compares the classical cloud/HPC schedulers and the proposed quantum scheduling framework, that highlight the key differences that necessitate a quantum-aware design.

The proposed framework should be capable of handling several failure modes that are unique to quantum systems. Since quantum states cannot be checkpointed, a job on a QPU that went offline must be executed on another QPU. Stale calibration data might cause the scheduler to assign a job to a degraded QPU, thus resulting in output below the required fidelity. The QoS feedback loop tries to mitigate this by triggering recalibration queries with the drop in result quality. The QPU's coherence time may be exceeded by a transpiled circuit, thus silently producing noisy output. The preprocessing service guards against this by comparing estimated execution time to coherence limits before scheduling. Under a heavy load, high-fidelity QPUs can become bottlenecks. By routing jobs that do not require top-tier fidelity to lower-performing but available QPUs, the best-fit fidelity policy alleviates this. Several considerations that are applied in classical scheduling and resource management cannot be applied in a quantum setup, as of today's position of quantum research. For example, allocating jobs based on the speed-accuracy tradeoff required by a user is a common practice in classical systems. However, Nakajima et al. [73] state that the existence of a speed-accuracy tradeoff in quantum systems is still doubtful. Scheduling together the tasks of different intensiveness (data-,

Aspect	Classical Cloud/ HPC Schedulers (e.g., SLURM, Kubernetes)	Proposed Quantum Framework
Hardware awareness	Static specifications (CPU cores, memory, GPU type)	Dynamic calibration data (gate fidelities, coherence times, qubit topology)
Job preemption	Supported via checkpointing and migration	Not possible since quantum state collapses upon measurement
Scheduling inputs	CPU/memory/GPU requirements and priority	Qubit count, circuit depth, minimum fidelity, coherence budget
Error handling	Retry, checkpoint-restart, migration to a healthy node	Full re-execution required since quantum states can't be checkpointed and error mitigation is applied in post-processing
Compilation/Transpilation	Not required or architecture-independent	Required; circuit must be transpiled to target QPU's native gate set and topology
Resource heterogeneity	Largely homogeneous clusters or labeled node pools	Fundamentally heterogeneous (different qubit technologies, topologies, and noise profiles)

Table II: Qualitative comparison of classical cloud/HPC schedulers and the proposed quantum scheduling framework

memory-, and compute-intensive) is also a common practice to achieve higher efficiency in classical systems. However, the state of research towards quantum memory itself is still nascent [74]. As the research in quantum progresses, all these aspects could be considered in quantum scheduling systems as well. A recent study [75] showed that it is possible to use multiple quantum computers as one, thus enabling the division of a quantum task among multiple quantum systems, which is a key direction in terms of scheduling and resource management in quantum systems.

Limitations of the proposed framework. We explicitly acknowledge that the framework presented in this section is conceptual and architectural. The framework has not been implemented, prototyped, or empirically validated on any quantum cloud platform. The reported complexity bounds are analytical estimates, not measured runtimes. We have discussed the failure-modes in a qualitative fashion. The real-world performance, including end-to-end scheduling latency, the accuracy of fidelity-aware allocation, and recovery behavior under QPU recalibration or outage events, will depend on several factors like workload mix, QPU heterogeneity, and the maturity of cloud APIs exposed by quantum providers. Therefore, we present the framework as a structured starting point for the community rather than a validated production system. We also identify empirical evaluation on emerging quantum cloud platforms (such as IBM Quantum, AWS Braket, and Azure Quantum) as a future work.

VIII. CHALLENGES AND FUTURE DIRECTIONS

Despite the several benefits of QAI in MC systems, the field is plagued with multiple challenges that future research should focus on. Some critical limitations of QAI in MC systems are shown in Figure 7 and discussed below along with future directions.

A fundamental tension exists between the capabilities of QAI and MC requirements. MC systems demand

determinism, reproducibility, and certifiability, the properties that are challenging for noisy and probabilistic NISQ devices to provide. Quantum measurements are inherently stochastic and device noise introduces run-to-run variability on current hardware. Consequently, near-term QAI is best suited for non-safety-critical subtasks and subsystems within MC pipelines. A few examples of such tasks are offline optimization, training-time feature extraction, pre-mission planning. As of current state of QAI research, safety-critical real-time decisions should remain under deterministic classical control. Full integration of QAI into safety-critical decision loops awaits fault-tolerant quantum hardware capable of providing the reproducibility guarantees that MC certification demands.

Trainability limits - barren plateaus: Variational quantum algorithms suffer from the barren plateaus problem - exponentially vanishing gradients with the increase in system size. This impedes hyperparameter tuning and learning. For broad classes of randomized PQCs, the expected gradient tends to concentrate near zero as qubit count increases, which makes optimization intractable without any tailored initialization or structural bias. Recent research [76] also shows that noise induces barren plateaus, thus placing a fundamental limit on trainability on NISQ-based hardware. The MC pipelines that rely on PQC-based training, like adaptive control or classification, must adopt provably barren plateaus, avoiding designs and shallow ansätze (problem-informed) to maintain learnability.

Noise and error mitigation: Error mitigation techniques like zero-noise extrapolation and probabilistic error cancellation are able to reduce the bias without the need for full error correction. However, they incur runtime and sampling overheads that grow with circuit size and noise strength. Long-term reliability of MC systems requires fault-tolerant quantum computing. Hence, techniques like surface codes that

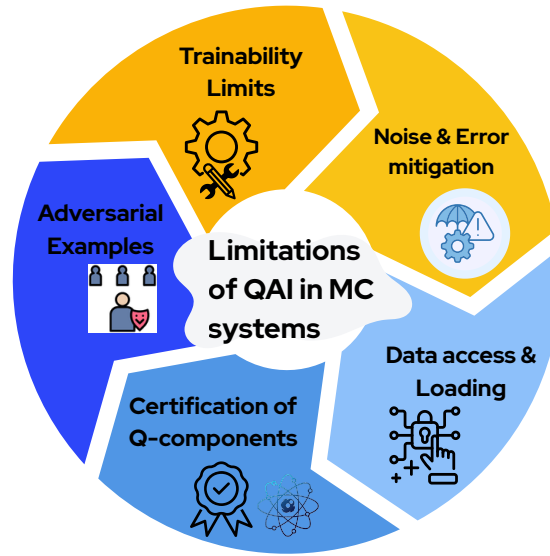


Figure 7: Limitations of Quantum AI in mission critical systems

quantify thresholds and the physical-qubit overheads to realize logical gates and qubits should be analyzed in MC systems. Error mitigation may just stabilize short-depth workloads in MC systems.

Data access and loading bottlenecks: There is a presumption of access to large classical datasets in quantum superposition in many stated QAI advantages. Quantum random-access memory (QRAM) architecture describes how 2^n cells can be accessed with n address qubits with only $O(\log N)$ active switches per query. However, this is a very non-trivial hardware assumption that highlights the data-loading problems as a distinct systems challenge [77]. Therefore, the MC systems with significant classical telemetry (like sensor fusion) must favor QAI formulations or engineer efficient state-preparation pipelines with less data-loading demands.

Certification and verification of quantum components: Before the quantum model can be deployed in the real-world mission critical systems, it is important to check that the model running in production is the same as validated in the lab. Hence, research must be directed towards developing tools for circuit equivalence checking that perform an automated check and prove if two circuits implement the same transformation after their compilation. QCEC (Quantum Circuit Equivalence Checking) tool [14] provides a good starting point in this direction. On the other hand, model-level verification frameworks need to be developed, which enable testing the robustness of a quantum classifier with a particular given input. Lin et al. [34] provided VeriQR in this direction.

Adversarial examples and robustness of QAI models: Recent research revealed that the quantum machine learning models can be fooled by tiny and carefully crafted input changes (adversarial examples), analogous to the classical machine learning models. Deploying such weak QAI models

in MC systems can have severe consequences. Although the research has progressed in the lines of quantification of robustness of QAI models [78], it is far beyond accurate adversarial robustness. The QAI models must therefore be treated as regular security-critical software, with routine stress and adversarial attack scenario tests. This is required until foolproof QAI models are developed.

The uncertainty of MC systems can be potentially handled by an alternative paradigm- fuzzy logic, which processes graded truth values through if-then rules that are human interpretable. Fuzzy logic offers a potential advantage for MC deployment in terms of its interpretability. Gentil et al. [13] proposed chemical pathways for fuzzy quantum computing. However, fuzzy systems require design of membership functions and rule bases by experts, they scale very poorly to high-dimensional input spaces, and they don't have the ability to learn representations from raw data. At the cost of stochasticity, limited interpretability, and certification challenges on current hardware, QAI can automatically extract features from high-dimensional sensor data and discover complex non-linear patterns. Thus, the two approaches are potentially complementary in nature, with fuzzy logic being better suited for subsystems with well-codified domain knowledge, while QAI might offer advantages for data-rich subtasks like anomaly detection and pattern recognition, which have already seen the deployment of classical ML.

IX. CONCLUSIONS

Quantum Artificial Intelligence has a significant potential to redefine the operational foundations of mission critical systems by combining the intrinsic parallelism and non-classical computation of quantum computing with the probabilistic reasoning of AI. In this paper, we examined how QAI

offers pathways to potentially address specific uncertainty management and computational bottlenecks that limit classical ML models across MC domains, ranging from autonomous defense networks to energy grids. However, several fronts have to be addressed before realizing the potential of QAI in MC systems, including the lack of standardized QAI benchmarks, the fragility of current-day quantum hardware, and the interpretability gap between the certifiable decision frameworks and quantum models. Future developments in QAI for MC systems are likely to emerge from hybrid quantum-classical models optimized for near-term NISQ devices that integrate quantum error correction tailored for AI workloads. A significant research direction in this field is the evolution of explainable QAI and quantum uncertainty to the level of verifiable standards in order to gain acceptance in safety-critical environments. Finally, the convergence of intelligent decision systems and quantum computing represents a paradigm shift towards intelligent systems that are secure, self-adaptive, and reliable in the face of uncertainty. It is also important to acknowledge that QAI may remain impractical for production MC deployment for years due to - limitations of NISQ hardware, the gap between QAI capabilities and MC certification requirements, and the pace of quantum hardware maturation. The value of the present review lies in mapping the research landscape so that progress toward this long-term goal can be structured and measured.

REFERENCES

- [1] A. Avizienis, J.-C. Laprie, B. Randell, and C. Landwehr, "Basic concepts and taxonomy of dependable and secure computing," *IEEE Transactions on Dependable and Secure Computing*, vol. 1, no. 1, pp. 11–33, 2004.
- [2] A. Chakraborty, M. Alam, V. Dey, A. Chattopadhyay, and D. Mukhopadhyay, "A survey on adversarial attacks and defences," *CAAI Transactions on Intelligence Technology*, vol. 6, no. 1, pp. 25–45, 2021.
- [3] G. Acampora, A. Chiatto, R. Schiattarella, and A. Vitiello, "Quantum artificial intelligence: A survey," *Computer Science Review*, vol. 59, p. 100807, 2026.
- [4] M. Baiocchi, F. Fagiolo, C. Loglisci, V. N. Losavio, A. Oddi, R. Rasconi, and P. L. Gentili, "Quantum artificial intelligence: Some strategies and perspectives," *AI*, vol. 6, no. 8, p. 175, 2025.
- [5] D. P. DiVincenzo, "The physical implementation of quantum computation," *Fortschritte der Physik: Progress of Physics*, vol. 48, no. 9–11, pp. 771–783, 2000.
- [6] D. Aharonov, M. Ben-Or, E. Eban, and U. Mahadev, "Interactive proofs for quantum computations," *arXiv:1704.04487*, 2017.
- [7] A. G. Fowler, M. Mariantoni, J. M. Martinis, and A. N. Cleland, "Surface codes: Towards practical large-scale quantum computation," *Physical Review A—Atomic, Molecular, and Optical Physics*, vol. 86, no. 3, p. 032324, 2012.
- [8] M. Kjaergaard, M. E. Schwartz, J. Braumüller, P. Krantz, J. I.-J. Wang, S. Gustavsson, and W. D. Oliver, "Superconducting qubits: Current state of play," *Annual Review of Condensed Matter Physics*, vol. 11, no. 1, pp. 369–395, 2020.
- [9] C. D. Bruzewicz, J. Chiaverini, R. McConnell, and J. M. Sage, "Trapped-ion quantum computing: Progress and challenges," *Applied Physics Reviews*, vol. 6, no. 2, 2019.
- [10] M. Saffman, "Quantum computing with atomic qubits and rydberg interactions: progress and challenges," *Journal of Physics B: Atomic, Molecular and Optical Physics*, vol. 49, no. 20, p. 202001, 2016.
- [11] F. Flamini, N. Spagnolo, and F. Sciarrino, "Photonic quantum information processing: a review," *Reports on Progress in Physics*, vol. 82, no. 1, p. 016001, 2018.
- [12] F. A. Zwanenburg, A. S. Dzurak, A. Morello, M. Y. Simmons, L. C. Hollenberg, G. Klimeck, S. Rogge, S. N. Coppersmith, and M. A. Eriksson, "Silicon quantum electronics," *Reviews of Modern Physics*, vol. 85, no. 3, pp. 961–1019, 2013.
- [13] P. L. Gentili, "Exploring new chemical paths in quantum ai: Preliminary accomplishments and future perspectives," *Advanced Quantum Technologies*, p. e00602, 2025.
- [14] L. Burgholzer and R. Wille, "QCEC: A JKQ tool for quantum circuit equivalence checking," *Software Impacts*, vol. 7, p. 100051, 2021.
- [15] J. Biamonte, P. Wittek, N. Pancotti, P. Rebentrost, N. Wiebe, and S. Lloyd, "Quantum machine learning," *Nature*, vol. 549, no. 7671, pp. 195–202, 2017.
- [16] M. Schuld and N. Killoran, "Quantum machine learning in feature hilbert spaces," *Physical Review Letters*, vol. 122, no. 4, p. 040504, 2019.
- [17] V. Havlíček, A. D. Córcoles, K. Temme, A. W. Harrow, A. Kandala, J. M. Chow, and J. M. Gambetta, "Supervised learning with quantum-enhanced feature spaces," *Nature*, vol. 567, no. 7747, pp. 209–212, 2019.
- [18] H.-Y. Huang, M. Broughton, J. Cotler, S. Chen, J. Li, M. Mohseni, H. Neven, R. Babbush, R. Kueng, J. Preskill *et al.*, "Quantum advantage in learning from experiments," *Science*, vol. 376, no. 6598, pp. 1182–1186, 2022.
- [19] M. Cerezo, A. Arrasmith, R. Babbush, S. C. Benjamin, S. Endo, K. Fujii, J. R. McClean, K. Mitarai, X. Yuan, L. Cincio *et al.*, "Variational quantum algorithms," *Nature Reviews Physics*, vol. 3, no. 9, pp. 625–644, 2021.
- [20] S. Sai, I. Goyal, V. Chamola, and R. Buyya, "Generative ai in the age of quantum computing: A taxonomy, architectural elements and future directions," *Authorea Preprints*, 2025.
- [21] S. Lloyd and C. Weedbrook, "Quantum generative adversarial learning," *Physical Review Letters*, vol. 121, no. 4, p. 040502, 2018.
- [22] S. Jerbi, L. M. Trenkwalder, H. Poulsen Nautrup, H. J. Briegel, and V. Dunjko, "Quantum enhancements for deep reinforcement learning in large spaces," *PRX Quantum*, vol. 2, no. 1, p. 010328, 2021.
- [23] M. Schuld, V. Bergholm, C. Gogolin, J. Izaac, and N. Killoran, "Evaluating analytic gradients on quantum hardware," *Physical Review A*, vol. 99, no. 3, p. 032331, 2019.
- [24] I. Cong, S. Choi, and M. D. Lukin, "Quantum convolutional neural networks," *Nature Physics*, vol. 15, no. 12, pp. 1273–1278, 2019.
- [25] L. Banchi, J. Pereira, and S. Pirandola, "Generalization in quantum machine learning: A quantum information standpoint," *PRX Quantum*, vol. 2, no. 4, p. 040321, 2021.
- [26] K. Temme, S. Bravyi, and J. M. Gambetta, "Error mitigation for short-depth quantum circuits," *Physical Review Letters*, vol. 119, no. 18, p. 180509, 2017.
- [27] E. Tang, "A quantum-inspired classical algorithm for recommendation systems," in *Proceedings of the 51st annual ACM SIGACT symposium on theory of computing*, 2019, pp. 217–228.
- [28] B. Khanal and P. Rivas, "Evaluating the impact of noise on variational quantum circuits in nisq era devices," in *Proceedings of the 2023 Congress in Computer Science, Computer Engineering, & Applied Computing (CSCE)*. IEEE, 2023, pp. 1658–1664.
- [29] J. Stokes, J. Izaac, N. Killoran, and G. Carleo, "Quantum natural gradient," *Quantum*, vol. 4, p. 269, 2020.
- [30] A. Pérez-Salinas, A. Cervera-Lierta, E. Gil-Fuster, and J. I. Latorre, "Data re-uploading for a universal quantum classifier," *Quantum*, vol. 4, p. 226, 2020.
- [31] S. Park and O. Simeone, "Quantum conformal prediction for reliable uncertainty quantification in quantum machine learning," *IEEE Transactions on Quantum Engineering*, vol. 5, pp. 1–24, 2023.
- [32] Q. Li, Y. Huang, X. Hou, Y. Li, X. Wang, and A. Bayat, "Ensemble-learning error mitigation for variational quantum shallow-circuit classifiers," *Physical Review Research*, vol. 6, no. 1, p. 013027, 2024.
- [33] J.-C. Huang, Y.-L. Tsai, C.-H. H. Yang, C.-F. Su, C.-M. Yu, P.-Y. Chen, and S.-Y. Kuo, "Certified robustness of quantum classifiers against adversarial examples through quantum noise," in *Proceedings of the IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP 2023)*. IEEE, pp. 1–5, 2023.
- [34] Y. Lin, J. Guan, W. Fang, M. Ying, and Z. Su, "Veriqr: A robustness verification tool for quantum machine learning models," *arXiv:2407.13533*, 2024.
- [35] M. C. Caro, M. Hinsche, M. Ioannou, A. Nietner, and R. Sweke, "Classical verification of quantum learning," *arXiv:2306.04843*, 2023.
- [36] K. Kadian, S. Garhwal, and A. Kumar, "Exqual: an explainable quantum machine learning classifier," *Applied Intelligence*, vol. 55, no. 13, pp. 1–21, 2025.
- [37] R. Heese, T. Gerlach, S. Mücke, S. Müller, M. Jakobs, and N. Piatkowski, "Explaining quantum circuits with shapley values:

- Towards explainable quantum machine learning,” *Quantum Machine Intelligence*, vol. 7, no. 1, pp. 1–33, 2025.
- [38] A. Kottahachchi Kankaname Don and I. Khalil, “Qrlaxai: quantum representation learning and explainable ai,” *Quantum Machine Intelligence*, vol. 7, no. 1, pp. 1–22, 2025.
- [39] L. Pira and C. Ferrie, “On the interpretability of quantum neural networks,” *Quantum Machine Intelligence*, vol. 6, no. 2, p. 52, 2024.
- [40] T. Stollenwerk, B. O’Gorman, D. Venturelli, S. Mandra, O. Rodionova, H. Ng, B. Sridhar, E. G. Rieffel, and R. Biswas, “Quantum annealing applied to de-conflicting optimal trajectories for air traffic management,” *IEEE Transactions on Intelligent Transportation Systems*, vol. 21, no. 1, pp. 285–297, 2019.
- [41] F. De Grossi, A. Carbone, D. Spiller, D. Ottaviani, R. Mengoni, and C. Circi, “Transcription and optimization of an interplanetary trajectory through quantum annealing,” *Astrodynamics*, pp. 1–21, 2025.
- [42] H. Makhanov, K. Setia, J. Liu, V. Gomez-Gonzalez, and G. Jenaro-Rabadan, “Quantum computing applications for flight trajectory optimization,” in *2024 International Conference on Quantum Communications, Networking, and Computing (QCNC)*. IEEE, pp. 65–74, 2024.
- [43] J.-M. Liang, S.-Q. Shen, M. Li, and L. Li, “Quantum anomaly detection with density estimation and multivariate gaussian distribution,” *Physical Review A*, vol. 99, no. 5, p. 052310, 2019.
- [44] W. Chen, J. Wan, F. Ye, R. Wang, and C. Xu, “Qmarl: A quantum multi-agent reinforcement learning framework for swarm robots navigation,” in *2024 IEEE International Conference on Acoustics, Speech, and Signal Processing Workshops (ICASSPW)*. IEEE, 2024, pp. 388–392.
- [45] J. Sun, Y. Meng, J. Huang, F. Liu, and S. Li, “Distributed cooperative control with collision avoidance for spacecraft swarm reconfiguration via reinforcement learning,” *Acta Astronautica*, vol. 205, pp. 95–109, 2023.
- [46] L. Evers, A. I. Barros, H. Monsuur, and A. Wagelmans, “Online stochastic uav mission planning with time windows and time-sensitive targets,” *European Journal of Operational Research*, vol. 238, no. 1, pp. 348–362, 2014.
- [47] N. Mori and S. Furukawa, “Quantum annealing for the adjuster routing problem,” *Frontiers in Physics*, vol. 11, p. 1129594, 2023.
- [48] A. Sinha, A. Macaluso, and M. Klusch, “Nav-q: quantum deep reinforcement learning for collision-free navigation of self-driving cars,” *Quantum Machine Intelligence*, vol. 7, no. 1, p. 19, 2025.
- [49] V. Heyraud, Z. Li, Z. Denis, A. Le Boité, and C. Ciuti, “Noisy quantum kernel machines,” *Physical Review A*, vol. 106, no. 5, p. 052421, 2022.
- [50] S. A. Khowaja, K. Dev, M. S. Pathan, E. Zeydan, and M. Debbah, “Integration of agentic ai with 6g networks for mission-critical applications: Use-case and challenges,” *arXiv preprint arXiv:2502.13476*, 2025.
- [51] S. Sai, I. Goyal, S. Sharma, S. H. Manuri, V. Chamola, and R. Buyya, “Quantum machine learning for cybersecurity: A taxonomy and future directions,” *arXiv:2512.15286*, 2025.
- [52] M. Alshowkan, P. G. Evans, M. Starke, D. Earl, and N. A. Peters, “Authentication of smart grid communications using quantum key distribution,” *Scientific Reports*, vol. 12, no. 1, p. 12731, 2022.
- [53] R. Kumar and M. Swarnkar, “Quids: A quantum support vector machine-based intrusion detection system for iot networks,” *Journal of Network and Computer Applications*, vol. 234, p. 104072, 2025.
- [54] M. Kalinin and V. Krundyshev, “Security intrusion detection using quantum machine learning techniques,” *Journal of Computer Virology and Hacking Techniques*, vol. 19, no. 1, pp. 125–136, 2023.
- [55] W. Li and D.-L. Deng, “Quantum delegated and federated learning via quantum homomorphic encryption,” *Research Directions: Quantum Technologies*, vol. 3, p. 3, 2025.
- [56] S. Sai, A. Sawaika, P. Singh, and R. Buyya, “Quantum federated learning: Architectural elements and future directions,” *arXiv:2510.17642*, 2025.
- [57] D. Bienstock and A. Verma, “Strong np-hardness of ac power flows feasibility,” *Operations Research Letters*, vol. 47, no. 6, pp. 494–501, 2019.
- [58] “What the duck curve tells us about managing a green grid,” [accessed 2025-10-17]. [Online]. Available: https://www.caiso.com/documents/flexibleresourceshelfrenewables_fastfacts.pdf
- [59] B. M. Peter and M. Korkali, “Quantum-enhanced reinforcement learning for power grid security assessment,” *arXiv:2504.14412*, 2025.
- [60] B. Salgado, A. Sequeira, and L. P. Santos, “A hybrid classical-quantum approach to highly constrained unit commitment problems,” *arXiv:2412.11312*, 2024.
- [61] A. Ajagekar and F. You, “Quantum computing based hybrid deep learning for fault diagnosis in electrical power systems,” *Applied Energy*, vol. 303, p. 117628, 2021.
- [62] K. A. Islam, D. Q. Chen, M. Marathe, H. Mortveit, S. Swarup, and A. Vullikanti, “Simulation-assisted optimization for large-scale evacuation planning with congestion-dependent delays,” *arXiv:2209.01535*, 2022.
- [63] Y. Liu, K. Komatsu, M. Kumagai, M. Sato, and H. Kobayashi, “Performance evaluation of tsunami evacuation route planning on multiple annealing machines,” in *Proceedings of the 20th ACM International Conference on Computing Frontiers*, pp. 185–188, 2023.
- [64] W. Fu, H. Xie, H. Zhu, H. Wang, L. Jiang, C. Chen, and Z. Bie, “Coordinated post-disaster restoration for resilient urban distribution systems: A hybrid quantum-classical approach,” *Energy*, vol. 284, p. 129314, 2023.
- [65] V. Sankaradass, M. Tholkapiyan, S. Sudhakar, and R. Devasenan, “Leveraging quantum machine learning for early warning systems in sudden environmental disaster prediction,” *Quantum Information Processing*, vol. 24, no. 9, pp. 1–14, 2025.
- [66] H. Kobayashi, “R&D of a quantum-annealing assisted next generation hpc infrastructure and its killer applications,” in *Proceedings of the Joint Workshops on Sustained Simulation Performance 2018 and 2019*. Springer, pp. 3–12, 2020.
- [67] “Sandboxaq completes major aqnav milestones with the usaf | sandboxaq,” 4 2025, [accessed 2025-10-21]. [Online]. Available: <https://www.sandboxaq.com/post/sandboxaq-completes-major-aqnav-milestones-with-the-usaf>
- [68] “Quantum technologies | airbus,” 12 2024, [accessed 2025-10-21]. [Online]. Available: <https://www.airbus.com/en/innovation/digital-transformation/quantum-technologies>
- [69] “News | ibm and raytheon technologies to collaborate on artificial intelligence, cryptography and quantum technologies | rtx,” [accessed 2025-10-21]. [Online]. Available: <https://www.rtx.com/news/2021/10/11/ibm-raytheon-technologies-to-collaborate>
- [70] “Enhancing solar power forecasting with hybrid quantum models: New study on cutting-edge methods,” [accessed 2025-10-21]. [Online]. Available: <https://terraquantum.swiss/news/new-study-by-terra-quantum-shows-quantum-models-outperform-classical-methods-for-solar-power-forecasting/>
- [71] “Quantum quants and tno:electrical grid optimization powered by quantum computing,” [accessed 2025-10-21]. [Online]. Available: https://www.dwavequantum.com/media/5q313uno/tnoqq_electrical-grid-optimization_case_study_2-0.pdf
- [72] “Iionq and hyundai motor company expand quantum computing partnership, continuing pursuit of automotive innovation,” [accessed 2025-10-21]. [Online]. Available: <https://iionq.com/news/iionq-and-hyundai-motor-company-expand-quantum-computing-partnership>
- [73] S. Nakajima and H. Tajima, “Speed-accuracy trade-off relations in quantum measurements and computations,” *arXiv preprint arXiv:2405.15291*, 2024.
- [74] R. Naik, N. Leung, S. Chakram, P. Groszkowski, Y. Lu, N. Earnest, D. McKay, J. Koch, and D. Schuster, “Random access quantum information processors,” *arXiv:1705.00579*, 2017.
- [75] A. Carrera Vazquez, C. Tornow, D. Riste, S. Woerner, M. Takita, and D. J. Egger, “Combining quantum processors with real-time classical communication,” *Nature*, vol. 636, no. 8041, pp. 75–79, 2024.
- [76] S. Wang, E. Fontana, M. Cerezo, K. Sharma, A. Sone, L. Cincio, and P. J. Coles, “Noise-induced barren plateaus in variational quantum algorithms,” *Nature communications*, vol. 12, no. 1, p. 6961, 2021.
- [77] V. Giovannetti, S. Lloyd, and L. Maccone, “Quantum random access memory,” *Physical Review Letters*, vol. 100, no. 16, p. 160501, 2008.
- [78] W. Gong and D.-L. Deng, “Universal adversarial examples and perturbations for quantum classifiers,” *National Science Review*, vol. 9, no. 6, 2022.

Manuscript Title: Quantum Artificial Intelligence for Mission-Critical Systems: Foundations, Architectural Elements, and Future Directions

Author Contributions:

Siva Sai — Data Curation, Methodology, Formal Analysis, Writing – Original Draft

Rajkumar Buyya — Conceptualization, Visualization, Writing – Review & Editing, Project Administration

All authors have read and approved the final version of the manuscript and agree to be accountable for all aspects of the work, ensuring that questions related to the accuracy or integrity of any part are appropriately investigated and resolved.

Declaration of interests

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

The authors declare the following financial interests/personal relationships which may be considered as potential competing interests:

Journal Pre-proof