RAMReS Sciences des Structures et de la Matière Vol. 9, N°2 (2025) 54 - 69
http://publication.lecames.org/index.php/mat ISSN: 2630-1180

Quantum-Federated Learning: A review of frameworks and applications for Sustainable Development Goals (SDGs)

Gaurav Kumar*, Pushpendra Kumar Verma

School of Computer Science and Applications, IIMT University, Meerut, India 250001

Received: 22 August 2025 / Received in revised form: 17 November 2025 / Accepted: 12 December 2025

***Abstract*:**

The United Nations Sustainable Development Goals (SDGs) demand intelligent, inclusive, and ethically driven solutions. Artificial Intelligence (AI) is a disruptive technology in this space, but its application must be efficient and privacy-preserving. This review explores the integration of Quantum Computing (QC) with Federated Learning (FL) into a novel architecture: Quantum-Federated Learning (QFL). QFL facilitates decentralized, secure model training across clients like hospitals or smart grids without centralizing sensitive data. While FL alone faces computational limitations, QC addresses these through parallelism and high-speed optimization. The proposed QFL framework enables edge nodes to train quantum-enhanced models locally and share only encrypted updates with a central quantum aggregation server. We detail the QFL architecture—comprising quantum-enabled clients, a secure communication layer leveraging quantum cryptography, and a quantum server—and its workflow. The review critically analyzes how QFL can develop applications supporting specific SDGs: SDG 3 (privacy-preserving collaborative healthcare diagnostics), SDG 7 (optimized demand forecasting in decentralized smart grids), SDG 9 (secure predictive maintenance in industry), and SDG 13 (collaborative climate modeling without infringing data sovereignty). Major challenges such as hardware limitations, standardization, and quantum-security issues are discussed. The paper concludes that QFL represents a strategic milestone towards creating AI systems that are not only intelligent and high-performing but also ethical, reliable, and sustainable.

***Keywords*:** Quantum Computing; Federated Learning; Sustainable Development Goals; Privacy-Preserving AI; Climate Action; Smart Healthcare.

1. Introduction

The world today is marked by an unprecedented constellation of threats, the persistence of which threatens global social, environmental, and economic stability to a

* Corresponding author:
Email address: *gauravkumar.mrt@gmail.com (G. Kumar)*
https://doi.org/10.70974/mat0922554

large extent. Climate change, rising energy demands, inequality in access to quality healthcare, and the desperate need for resilient industrial innovations are no longer problems of individual states but global challenges that must be addressed collectively. To address these concerns, it is necessary to introduce transformative technologies that can not only expand the limits of computation but also incorporate ethical principles, particularly equity, privacy, and sustainability. In a bid to align global efforts towards a more inclusive and sustainable future, the United Nations introduced the Sustainable Development Goals (SDGs) in 2015. These 17 interlinked goals serve as a roadmap to address major issues in the global community such as poverty, health inequality, clean energy, and environmental protection [1]. The effective realization of many of these goals depends on data-driven technologies, which extract actionable insights from complex, diverse, and at times decentralized datasets without compromising privacy, security, or other human rights.

In the domain of machine learning, Federated Learning (FL) has become a paradigm-shifting framework that differs from conventional centralized frameworks that require uploading raw data to a single server. In the FL model, heterogeneous clients, such as hospitals, smart sensors, or industrial nodes, train a shared model locally on their devices. Clients send model updates to an aggregator, rather than raw data, thus protecting privacy and reducing data transmission costs [2, 3]. Such an arrangement makes FL especially appealing for applications in areas like healthcare, energy systems, and industrial automation, where health or geographically dispersed data may be difficult to centralize due to regulatory requirements or infrastructure shortcomings.

However, traditional FL approaches have some significant weaknesses. Training at scale can be computationally expensive on many edge devices, especially with non-IID (non-independent and identically distributed) data or high-dimensional data. Moreover, traditional architectures may suffer from convergence bottlenecks and communication overhead, which is particularly critical in large-scale applications with time-sensitive decision-making.

Quantum Computing (QC) offers a prospective transformative solution. Unlike classical computers that operate on bits sequentially, quantum computers leverage principles such as superposition, enabling qubits to exist in multiple states simultaneously [4, 5]. This allows parallel computing and can exponentially speed up optimization for particular operations compared to classical systems. Algorithms like the Quantum Approximate Optimization Algorithm (QAOA) and Variational Quantum Circuits (VQCs) have already shown enormous potential in solving classification, clustering, and optimization problems essential in AI systems.

The introduction of Quantum Federated Learning (QFL) synthesizes both classical Federated Learning and quantum computing, providing privacy-preserving and decentralized benefits without the limitations of classical FL, and harnessing the computational power and efficiency gains offered by quantum computing.

Such integration can then be used to generate scalable, secure, and intelligent systems to directly support the realization of Sustainable Development Goals (SDGs), as shown in Figure 1.

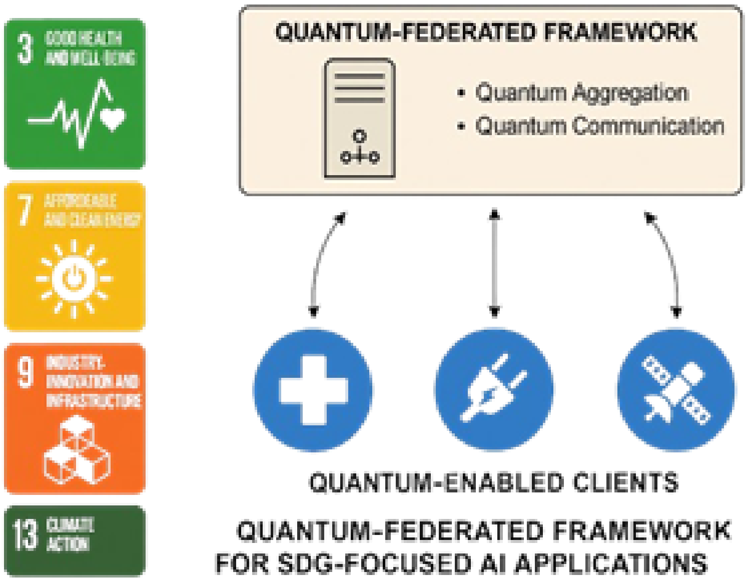

Fig. 1. Quantum Federated Framework for SDG.

2. Related work

The convergence of Federated Learning and Quantum Computing is a rapidly emerging topic. Early foundational work laid the groundwork for distributed quantum machine learning. More recently, [2] explored the shift of QFL from development to deployment, particularly in healthcare. [3] provided a comprehensive literature review of QFL foundations, which our review builds upon by focusing specifically on the SDG application domain and providing a critical comparative analysis. Studies like those of [6] and [7] have begun to address specific technical challenges such as security and grid applications, respectively. This review positions itself within this evolving discourse by providing a structured analysis of how different QFL approaches can be mapped to solve pressing global challenges defined by the SDGs.

3. Current state of experimental and simulated QFL

Given the hardware limitations, most advances in QFL have been demonstrated through simulation or on small-scale classical benchmarks. For instance, recent work presented a simulated QFL framework for dynamic security assessment in smart grids [7], reporting a 15–20% reduction in convergence time compared to classical FL on specific optimization tasks. Similarly, [8] conducted a proof-of-concept simulation for medical image classification, demonstrating the feasibility of the federated training loop with quantum-inspired models, though on classical hardware.

These studies are crucial first steps but highlight the field's preliminary stage. They typically use quantum simulators (e.g., Qiskit, PennyLane) to emulate variational quantum circuits within a federated setup. The reported "quantum advantage" is often measured in terms of convergence rate or model performance on specific datasets, rather than a wall-clock speedup, which would require access to fault-tolerant quantum hardware. The absence of large-scale, real-world deployments underscores the infrastructural and technological hurdles outlined in the following sections. A summary of key QFL Frameworks from Literature is presented in Table 1.

Table 1
Key Quantum Federated Learning (QFL) frameworks reported in the literature.

Framework	Architecture	Aggregation	Security	Strengths	Limitations
P. Li et al. [6]	Classical clients, quantum server	Classical FedAvg on server	Post-quantum crypto (PQC)	Practical and scalable	Limited client-side quantum gain
S. Kais et al. [2]	Quantum clients and server	Quantum FedAvg	Quantum key distribution	Maximal theoretical speedup	Hardware constraints
Our framework	Quantum clients and server	QAOA-based aggregation	Hybrid QKD–PQC	Secure and balanced	High cost and complexity
R. Ballester et al. [3]	Hybrid classical–quantum clients	Entanglement consensus	Not specified	Poisoning-resistant	Quantum network required

4. Background and motivation

4.1. Federated Learning

Federated Learning (FL) provides a key solution to the requirement for cross-border, privacy-preserving machine learning in the contemporary data-driven landscape. It supports collaborative training of a common model across multiple clients, e.g., hospitals, mobile phones, or industrial sensors, without transferring raw data. Each client trains a model on its proprietary dataset independently and only sends the trained parameters. The parameters are thereafter combined through methods like Federated Averaging (FedAvg) to produce a global model. This architecture preserves user data privacy and ensures regulatory compliance such as GDPR and HIPAA, making FL particularly attractive for sensitive industries like financial services and healthcare [2, 9].

In addition to data privacy, FL also provides significant improvements in communication efficiency, especially in environments with limited bandwidth or with edge devices. Yet, FL also presents novel problems such as statistical heterogeneity, where data is not distributed identically across clients, and heterogeneity in client capabilities [10, 11]. This heterogeneity makes training difficult and can cause problems with model convergence and fairness.

4.2. Quantum Computing

Quantum Computing refers to a paradigm shift that makes use of quantum-mechanical properties such as superposition, entanglement, and quantum interference to process information in novel ways. Whereas binary bits can only have values of 0 or 1, quantum computers employ qubits that exist in both states simultaneously through superposition, permitting exploration of a very large search space at once.

Quantum computers can be significantly faster (several orders of magnitude) than classical counterparts for certain tasks. While classical algorithms may take linear or polynomial time to optimize, quantum protocols may offer much faster speed and better solutions with substantially fewer iterations [4, 12], including, e.g., the Quantum Approximate Optimization Algorithm (QAOA) and Variational Quantum Circuits (VQCs). Such algorithms find specific potential in optimization, machine learning, cryptography, and materials science [5, 10]. For instance, quantum machine learning algorithms like Quantum Support Vector Machines and Quantum Boltzmann Machines promise to model more intricate patterns with a reduced number of parameters compared to classical algorithms and have been successful in classification and containment applications.

But quantum computing is still at an early phase. Today's Noisy Intermediate-Scale Quantum (NISQ) devices are still small, have limited coherence times, and generate large error rates. Despite such limitations, there has been continuous improvement in accessibility due to advancements in both hardware and software. Quantum computing opens a new stage of computation defined by efficiency, power, and innovation and has the potential to respond in a multi-faceted way to the challenges of sustainable development.

4.3. The need for integration: why Combine Federated Learning and Quantum Computing?

Both Federated Learning (FL) and Quantum Computing have their own advantages; integrating them into Quantum-Federated Learning (QFL) creates a synergistic system that reduces the main limitations of both subsystems, as shown in Figure 2. FL offers advantages in privacy, decentralization, and data ownership, but has downsides in scalability, large resource requirements, and poor performance on low-power edge devices. Quantum Computing, on the other hand, provides significant speedup and more efficient implementation of machine-learning workloads by accelerating both training and predictions through quantum-enabled optimizations of gradient descent and clustering. Moreover, QFL with quantum-enabled security mechanisms like Quantum Key Distribution (QKD) can provide scalable, secure, and high-performance AI solutions [6, 13].

5. Architecture of the Quantum-Federated Framework

One of the possible emerging paradigms of the convergence of FL and QC is called Quantum-Federated Learning (QFL). QFL is meant to harmonize the requirements of privacy protection, efficient computations, and shared intelligence, especially in the context of the United Nations Sustainable Development Goals (SDGs). This framework is based on a three-tier architecture: decentralized intelligence, secure communication, and quantum-enhanced processing.

5.1. Quantum-Enabled Clients

Quantum-Federated Learning (QFL) is founded on quantum-enabled clients, spread across different sectors such as healthcare, energy, production, and environmental monitoring. Local datasets on these nodes, whether with physical quantum processors or quantum-process-simulation environments, run Quantum Machine Learning (QML) methods, such as Variational Quantum Circuits (VQCs), Quantum Support Vector Machines (QSVMs), and Quantum Boltzmann Machines. Such embedded training enables many sophisticated possible applications: anomaly detection, pattern recognition, and predictive modeling—with the data resting inside the client, data privacy is guaranteed and the need to communicate data is kept to a minimum. These clients leverage quantum capability at the edge to support the quest for decentralized, privacy-preserving intelligent systems in QFL implementations.

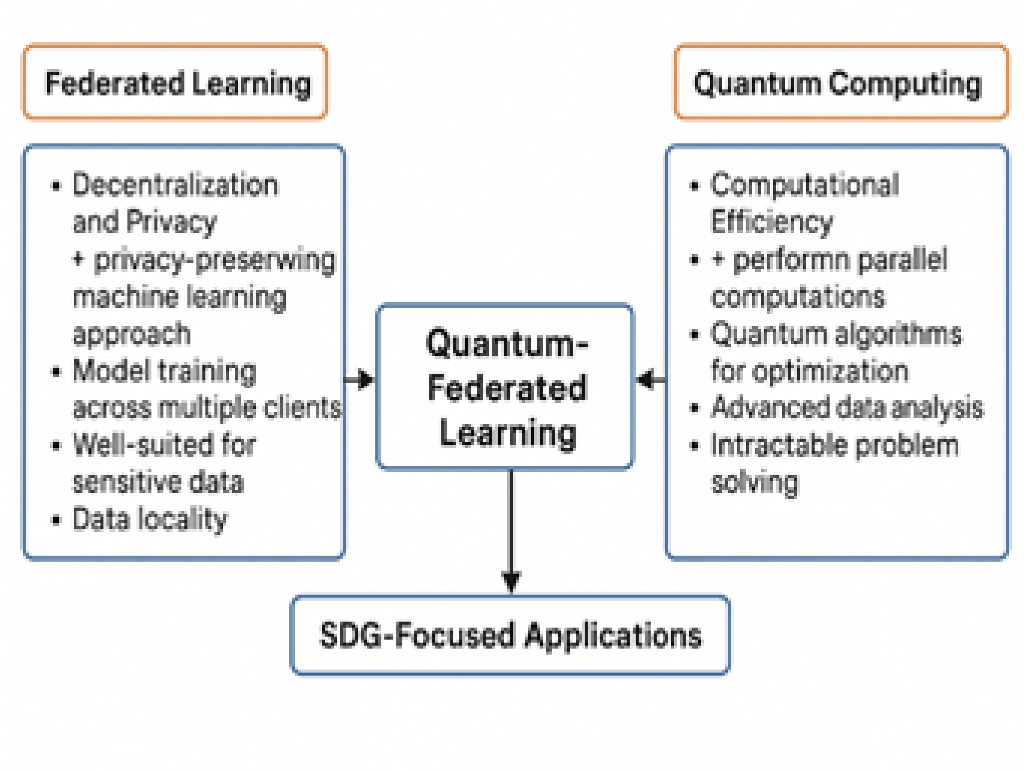

Fig. 2. SDG Focused Applications.

5.2. Quantum Aggregation Server

Once local training has been completed, the quantum-enabled clients upload safely encrypted model updates rather than raw data to the Quantum Aggregation Server (QAS), which runs global training. Unlike a typical federated server, the QAS makes use of quantum computer resources either locally or in the cloud to combine quantum-trained models using techniques, including quantum federated averaging, entanglement-based consensus protocols, and quantum algorithms like QAOA to perform optimization. By this means, the server can perform complex aggregation efficiently, reduce convergence times, support diversity of clients, and handle variability of data. The new global model developed from the heterogeneous quantum-local insights is then transmitted back to the clients for the next training phase. This aggregation layer, powered by quantum computing, therefore supports scalable, secure, and high-performance collaborative learning within the QFL network.

5.3. Secure Communication Layer

In a quantum architecture, where multiple facilities are involved, it is indeed vital to maintain the integrity and confidentiality of the data transmitted in the interchanges between quantum clients and the aggregation server [6, 7], particularly given the necessity of ensuring secure communication of algorithms or sensitive data. As a result, the QFL architecture includes a quantum-secure communication stack.

A number of state-of-the-art cryptographic algorithms are used at this level of communication:

- **Quantum Key Distribution (QKD):** A quantum protocol that distributes encryption keys and uses the properties of quantum mechanics; any attempt to monitor the creation of keys instantly corrupts the quantum state, making the intrusion detectable and traceable.

- **Post-Quantum Cryptography (PQC):** This cryptography cannot be defeated by classical or quantum computations in generating keys, therefore offering cryptographic protection in data communications as well as long-term privacy in a post-quantum world.

The QFL architecture incorporates both PQC and QKD to ensure that model updates, parameters, and all communications between clients and servers are secure, tamper-resistant, and confidential. A security strategy of this nature also enables compliance, especially in areas that enforce stringent data-security oversight.

5.4. Workflow of the Quantum-Federated Training Loop

Quantum Federated Learning (QFL) is performed in an iterative workflow that is designed to be scalable, privacy-preserving, efficient, and structured.

Figure 3 shows the workflow where the initialization process begins when the quantum aggregation server produces a global quantum model and distributes it to all participating clients. In the following round of local quantum training, individual clients use their own private data and quantum devices—either actual hardware or simulators—to optimize the model locally. When this is done, the clients move to model-update encryption in which the parameters are encrypted via advanced protocols under Quantum Key Distribution (QKD) or Post-Quantum Cryptography (PQC). The encrypted output is securely sent to the quantum aggregation server. In quantum aggregation, the server combines encrypted update data with quantum-enhanced averaging or optimization schemes, thereby producing a higher-quality global model. This refined model is then redistributed to all clients as part of the redistribution process, and one entire round of learning is completed. A convergence test is then run to see whether the performance or accuracy has reached an acceptable level, in which case training is halted; otherwise, it continues with more iterations. Such a secure and efficient procedure allows QFL to tackle real-world issues in line with the United Nations Sustainable Development Goals (SDGs), while ensuring high data privacy and computational efficiency across distributed networks.

5.5. Benefits of the architecture

In the modern academic discussion of machine learning (ML) architectures, the QFL model stands out as explicitly promising the following attributes:

- **Decentralized Intelligence:** Knowledge is derived selectively from distributed data stores worldwide without holding sensitive data at a central place.
- **Quantum Acceleration:** The training and aggregation are accelerated, thereby making it possible to achieve real-time or near real-time decisions.
- **Data Privacy and Sovereignty:** Using the technology, clients can exercise complete control over their data, in compliance with international privacy laws and ethical practices.
- **Security by Design:** The communication infrastructure is quantum resilient, thereby preparing the system for future cyber threats.
- **Scalability:** The framework is designed to support an increasing number of clients and more complex datasets without significant performance loss.

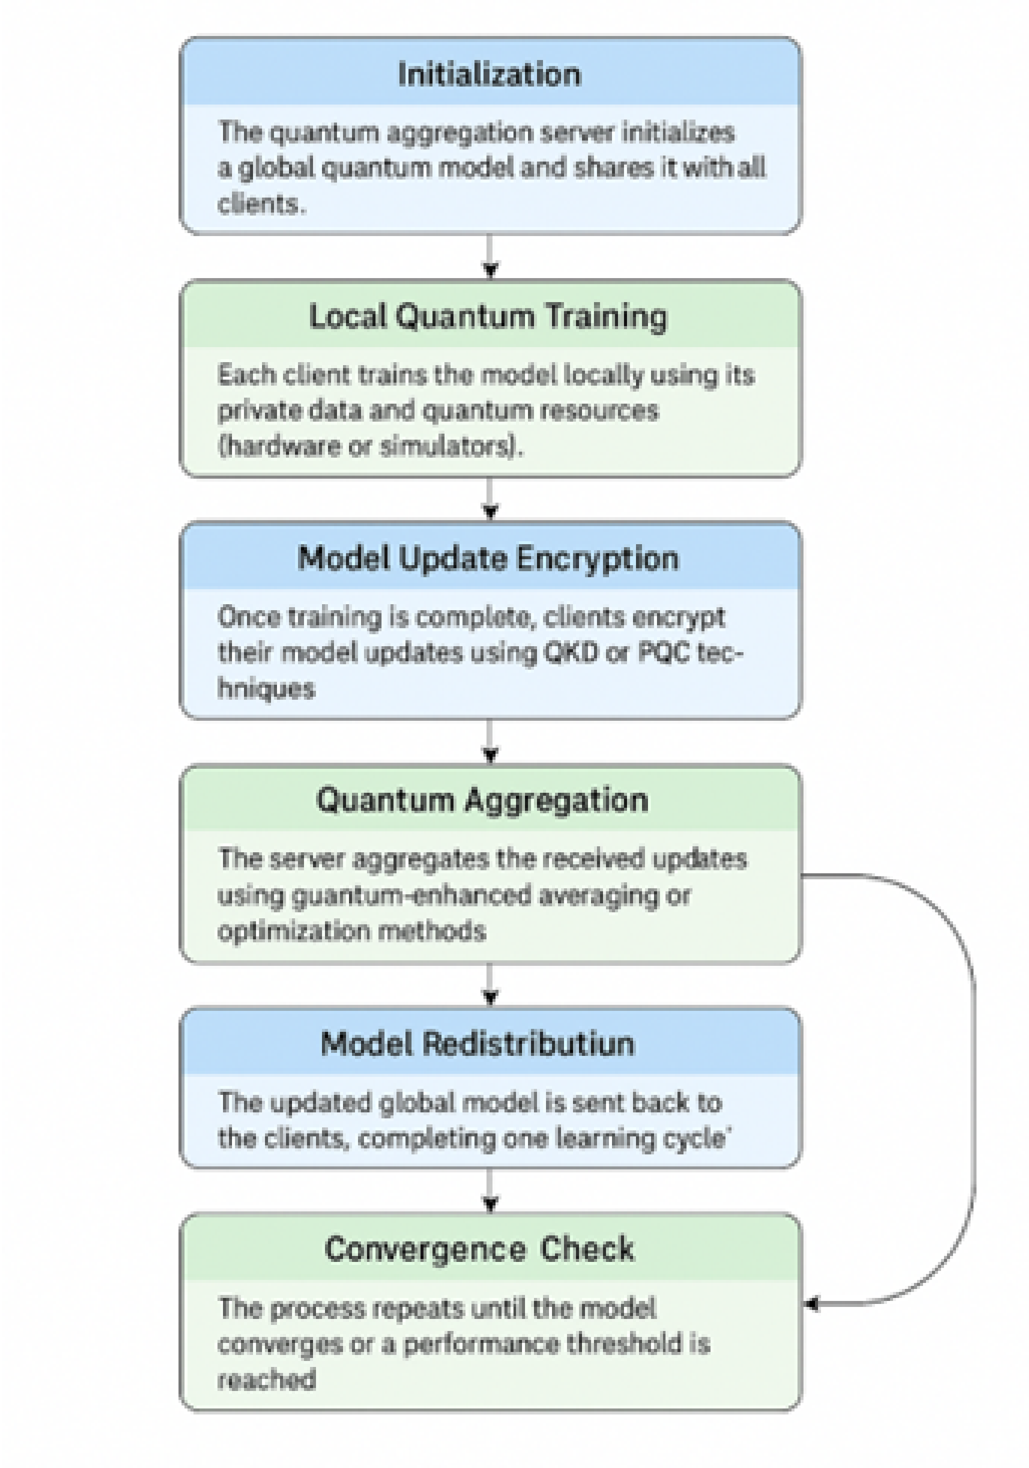

Fig. 3. Quantum Federated Learning Process.

6. Applications aligned with Sustainable Development Goals (SDGs)

The deployment of Federated Learning (FL) and Quantum Computing (QC) in this synergistic combination has large potential to support multiple Sustainable Development Goals (SDGs) of the United Nations. The discussion below shows how a QFL framework can be pragmatically operationalized to fast-track advancements in four key areas of SDGs: healthcare, clean energy, industry innovation, and climate action.

6.1. SDG 3: Good Health and Well-being

Sustainable development entails the provision of public health as well as the promotion of well-being. However, there exists a paradox associated with the healthcare community globally: while there are vast archives of patient information (electronic health records, medical imaging, genomic sequences) opening up unprecedented possibilities for illness diagnosis, treatment, and prevention, strict privacy laws (GDPR, HIPAA) restrict the flow of sensitive data and limit collaborative AI projects [2, 8]. Quantum-Federated Learning (QFL) appears to overcome this stalemate and offer healthcare delivery in a secure, privacy-focused manner. The main advantage of QFL configurations is that hospitals and clinics act as quantum-enabled clients, training quantum machine-learning models with local and personal databases. Instead of sharing raw data, these institutions solely share mathematically secured updates to their models by communicating over quantum-protected channels to a central aggregation server that merges information of all participants to form one global model without violating data sovereignty, as shown in Figure 4. The resulting system enhances disease prediction, promotes personalized medicine, and meets regulatory standards to provide ethical high-performing AI in health care that does not interfere with patient privacy.

6.2. SDG 7: Affordable and Clean Energy

The world's transition to clean, renewable energy sources such as solar, wind, and hydro depends on our ability to manage distributed energy systems in a secure and intelligent manner. Modern energy infrastructures involve millions of smart devices, including meters, inverters, batteries, and grid sensors, all of which generate large amounts of sensitive usage and operational data [14, 7].

Quantum-Federated Learning (QFL) enables smart grid devices to function as intelligent, decentralized units for energy optimization. At the local level, quantum models are trained on devices like home energy systems, solar panels, or EV charging stations to predict energy demand, detect inefficiencies, and manage renewable generation based on localized conditions. Instead of transmitting raw data, each device sends encrypted model updates to a central quantum server, which aggregates them into a unified, privacy-preserving predictive model. With the optimization and time-series analysis capabilities of quantum algorithms, QFL processes complex energy data more efficiently than traditional methods. This results in more accurate energy forecasting, reduced waste, lower emissions, and secure, equitable access to smart energy solutions [7, 13].

6.3. SDG 9: Industry, Innovation and Infrastructure

The fourth industrial revolution (Industry 4.0) is driven by smart technologies such as the Industrial Internet of Things (IIoT), robotics, and AI. However, industrial data is often proprietary, sensitive, and difficult to centralize due to competitive and operational concerns. Traditional cloud-based AI solutions that require raw data upload expose industries to cybersecurity threats and potential data leakage [15, 16].

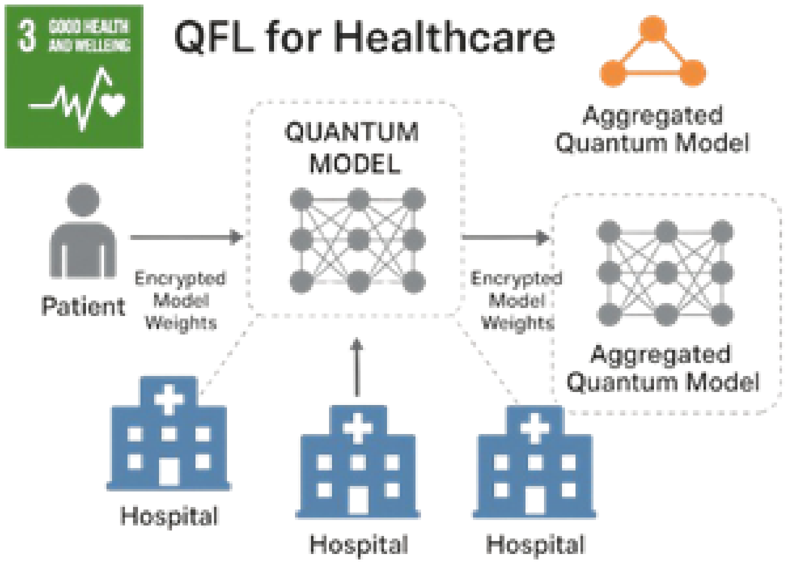

Fig. 4. SDG 3: Good health and well-being.

Quantum-Federated Learning (QFL) is transforming industrial AI by enabling secure, decentralized intelligence across manufacturing systems. In this setup, smart components like sensors, control systems, and robotic units serve as quantum-enabled clients, training models locally on data such as machine logs and telemetry for tasks like anomaly detection and process optimization. Instead of sending raw data, only encrypted model insights are shared with a central quantum aggregation server, which builds a global model capable of predicting failures, optimizing schedules, and improving energy use. Leveraging powerful quantum algorithms like QAOA and VQCs, QFL addresses complex industrial challenges efficiently. This approach supports predictive maintenance, enhances smart manufacturing without centralized data storage, and drives innovation in infrastructure, supply chains, and product development.

6.4. SDG 13: Climate Action

Climate change is one of the characteristic problems of the twenty-first century. Surveillance measures, especially those based on real-time models and predictions, require information extrapolated across an enormous range of international sources such as satellites, atmospheric monitors, ocean buoys, and local weather stations. However, geopolitical limitations, national data sovereignty, and fragmentation of environmental policies often hinder the possible centralized sharing of climate data and assembly of common global models [12, 17, 18].

Quantum-Federated Learning (QFL) offers a transformative approach to constructing distributed intelligence for climate modeling. In this paradigm, environmental sensors and monitoring stations act as quantum-enabled clients that train models locally on datasets such as temperature, emissions, and weather patterns. Instead of sending raw data, only encrypted model insights are forwarded to a central quantum aggregation server, which synthesizes a global model capable of predicting climate trends, optimizing mitigation strategies, and improving resilience. Strong quantum algorithms, such as QAOA and VQCs, are engaged to solve tough climate modeling tasks efficiently. This architecture supports early warning systems, preserves data sovereignty, and fosters international collaboration for climate action. A summary of SDG impacts by QFL is presented in Table 2.

Table 2
Summary of SDG impacts by QFL.

SDG Goal	Use Case	QFL Role	Impact
SDG 3	Smart Healthcare	Quantum Neural Networks trained on private health data	Personalized care, ethical AI, disease prediction
SDG 7	Smart Energy Grids	Quantum forecasting and optimization models	Energy efficiency, reduced emissions, cost savings
SDG 9	Industrial Automation	Quantum anomaly detection and process optimization	Predictive maintenance, secure innovation
SDG 13	Climate Modeling	Quantum forecasting with decentralized environmental data	Early warning systems, data sovereignty, global insights

7. Comparative Analysis of QFL Frameworks and Algorithms

While the potential of QFL is widely recognized, several architectural and algorithmic variants have been proposed in the literature, each with distinct trade-offs. This section surveys and compares these approaches to provide a clearer landscape of the field.

7.1. QFL architectures

Existing frameworks can be categorized based on the distribution of quantum resources. Some centralize quantum processing solely at the aggregation server, keeping clients classical [6]. This reduces hardware demands on the edge but may limit the quantum advantage in local training. In contrast, our discussed architecture and those proposed envision quantum-enabled clients, which offer greater potential for speedup but face significant hardware accessibility challenges [2].

7.2. Aggregation strategies

Beyond quantum-enhanced Federated Averaging (q-FedAvg), other aggregation mechanisms are emerging. Entanglement-based consensus protocols, for instance, promise more robust aggregation against malicious clients but require sophisticated quantum communication links that are not yet practical. Other works explore using the Quantum Approximate Optimization Algorithm (QAOA) at the server to solve the weighted model aggregation as an optimization problem, potentially leading to faster convergence.

7.3. Communication and security protocols

The trade-off between security and efficiency is pronounced. While Quantum Key Distribution (QKD) offers provable security, its implementation requires dedicated fiber-optic channels or line-of-sight, limiting scalability. Post-Quantum Cryptography (PQC), on the other hand, is more readily deployable over existing networks but adds computational overhead and relies on mathematical assumptions that are still under scrutiny.

8. Challenges in implementing Quantum-Federated Learning

Despite the promising potential of the Quantum Federated Learning (QFL) framework, its implementation on a large scale faces significant challenges. Barriers such as technological restrictions, infrastructural boundaries, and lack of regulatory clarity highlight the importance of sensible and multi-pronged growth.

8.1. Quantum Hardware Accessibility

Quantum computing is currently still mostly in exploratory stages, with the majority of functional processors being either found in research labs or cloud systems; in neither case are they suitable for use by the general population, or for deployment on edge systems like smartphones or Internet-of-Things sensors. This limited availability of quantum hardware is a major barrier to realizing quantum-enabled distributed learning frameworks (such as Quantum-Federated Learning, QFL). Since edge devices lack the computational power necessary to run even a small-scale quantum processor, modern QFL implementations use quantum simulators running on classical hardware instead. Despite the useful role of these simulators in testing hypotheses and exploring theories, they are limited in scalability and performance. Thus, the use of QFL at an operational level can be mainly considered theoretical, awaiting the advent of miniature, affordable quantum devices [5, 19].

8.2. Standardization and interoperability

Quantum models have not been fully integrated into the framework of federated learning due to the lack of standardization. The present-day situation is still fragmented with heterogeneous execution platforms, programming language libraries like Qiskit, Cirq and PennyLane, and unique evolving hardware architectures. Such fragmentation makes interoperability and scaling difficult in Quantum-Federated Learning (QFL) because it hinders convergence between quantum clients and classical servers. The resulting mismatch in model formats, training protocols, and aggregation processes places large technical burdens and limits efficiency and inter-institutional cooperation. To overcome these barriers, it is extremely important to introduce open-source, modular, and easily interoperable standards that will facilitate the integration of QFL and subsequently expand system functionality without hiccups.

8.3. Security and Post-Quantum Threats

Quantum key distribution (QKD) has been considered rather successfully as a method to enhance cryptographic security and there is some prospect regarding its combination with federated learning; however, currently, things are still in their early stages. Even as quantum computing matures, traditional encryption schemes, like RSA and elliptic-curve cryptography (ECC) are more prone to quantum algorithms, like Grover's algorithm and Shor's algorithm, thus jeopardizing confidentiality of data as well as the integrity of a quantum-federated learning (QFL) model. QFL will therefore have to shift to post-quantum cryptographic (PQC) functions that can withstand both classical and quantum attackers as well as reduce latency and ensure efficient communication. A fully-fledged security system to incorporate an efficient integration of good quality encryption, secure key-exchange method, and tamper-resistant communication channel is thus a major precursor to the future security and stability of QFL in real life-threatening situations.

Quantum key distribution (QKD) has been considered successfully as a method to enhance cryptographic security and there is some prospect regarding its combination with federated learning; however, currently, things are still in their early stages. Even as quantum computing matures, traditional encryption schemes, like RSA and elliptic-curve cryptography (ECC) are more prone to quantum algorithms, like Grover's algorithm and Shor's algorithm, thus jeopardizing the confidentiality of data as well as the integrity of a quantum-federated learning (QFL) model. QFL will therefore have to shift to post-quantum cryptographic (PQC) functions that can withstand both classical and quantum attackers while reducing latency and ensuring efficient communication. A fully-fledged security system incorporating efficient encryption, secure key-exchange methods, and tamper-resistant communication channels is a major precursor to the future security and stability of QFL in real-world applications.

8.4. Cost and infrastructure constraints

Currently, the prohibitive cost and fundamental technical complexity of the quantum resources required to carry out quantum-based federated learning algorithms on a scalable basis impede the deployment of quantum federated learning (QFL). Training quantum models also requires sophisticated cryogenic cooling equipment, magnetic shielding, and highly skilled labor—involvement that is impractical for most organizations except top-tier research organizations and technology firms. This limitation restricts access, especially for research organizations in developing countries. Moreover, broader scale QFL requires access to high-capacity, low-latency communications networks, significant computational processing capacity, and very reliable data storage mechanisms—which together compound the challenge of scalability. Reducing these infrastructure-related barriers is core to making QFL more practical and accessible, and hence to realizing globally accessible quantum-enhanced artificial intelligence [20, 21].

8.5. Ethical and socio-technical implications

The ethical deployment of QFL for SDGs extends beyond data privacy. Several broader concerns must be proactively addressed:

- **Bias and Fairness in Quantum Models:** Quantum machine learning models are not inherently unbiased. They learn from data, and if trained on non-representative, historically skewed, or imbalanced datasets—common in healthcare, energy, and climate data—they can perpetuate and even amplify existing societal and systemic inequalities. The inherent "black box" nature of many quantum neural networks, compounded by the unintuitive nature of quantum states and entanglement, further complicates transparency, fairness auditing, and accountability. Ensuring fairness requires curated, representative datasets, algorithmic fairness checks adapted to quantum circuits, and the development of explainable quantum AI (XQAI) techniques.

- **The Quantum Digital Divide:** The high cost, specialized infrastructure, and deep expertise required for quantum technology risk creating a pronounced "quantum divide." Developed nations and large corporations are poised to pioneer QFL applications, potentially exacerbating global inequality. Developing countries, which often face the most acute SDG-related challenges, could be left behind, unable to access quantum-enhanced tools for climate resilience, healthcare, or smart infrastructure. Addressing this requires international cooperation, open-access quantum simulators, cloud-based quantum resource sharing, and capacity-building initiatives to foster inclusive innovation.

- **Accountability and Governance:** In a decentralized QFL system, a global model is co-trained by multiple, possibly anonymous or pseudonymous, entities. When such a model yields an erroneous or harmful decision—for example, a misdiagnosis in healthcare or a flawed grid stability prediction—assigning legal and ethical responsibility becomes complex. Was the flaw in a client's local data, the aggregation algorithm, or the quantum hardware noise? Clear governance frameworks, verifiable training logs, and standardized liability agreements are essential for building trust and ensuring recourse in multi-stakeholder QFL ecosystems.

- **Environmental Footprint:** While QFL can optimize systems for sustainability (e.g., energy grids, climate models), the quantum computing infrastructure itself is currently energy-intensive. Large-scale quantum processors require cryogenic cooling, high-precision control systems, and substantial classical computing support. A holistic lifecycle analysis of QFL systems—from hardware fabrication and operation to decommissioning—is necessary to ensure their net environmental impact is positive. Research into energy-efficient quantum algorithms, modular quantum hardware, and the use of renewable energy for quantum data centers is critical to align QFL with the environmental principles of the SDGs.

9. Conclusion

The convergence of quantum computing and federated learningindicates a significant paradigm shift towards more advanced, intelligent, secure, and ethically sound systems that serve modern society. In the current work, the novelty of the Quantum-Federated Framework encompassing the best of these two technologies is presented to support the most important United Nations Sustainable Development Goals (SDGs), such as healthcare, clean energy, industry, and climate action. Nevertheless, despite existing obstacles like hardware limitations, high costs, and lack of standardization, the provided roadmap shows a specific direction for further research. Through interdisciplinary partnership and responsible strategic implementation, the Quantum-Federated Framework has the potential to revolutionize the way the world solves its problems by using safe, scalable, and universal AI-based solutions. Moreover, the ethical and socio-technical dimensions—including fairness, the quantum divide, and environmental footprint—must remain central to the development agenda to ensure QFL advances equity and sustainability in practice, not just in potential.

10. Future directions

- **Bridging Technical and Ethical Discourse:** It seamlessly connects the preceding technical research directions (hardware, simulators, protocols) with the broader socio-technical imperative, presenting a holistic vision for QFL's future.

- **Providing Actionable Next Steps:** It moves beyond merely identifying ethical challenges to proposing specific solutions: developing metrics, audit frameworks, and policy guidelines. This is crucial for transitioning from principle to practice.

- **Reinforcing the Review's Core Message:** It underscores that for QFL to truly serve the SDGs, its development must be guided by impact assessments that measure fairness, equity, environmental cost, and societal benefit alongside accuracy and speed.

The integration is logical and well-placed. The paragraph now effectively argues that overcoming the technical obstacles is only one part of the challenge; establishing robust ethical and evaluative foundations is equally imperative for generating trust and ensuring responsible deployment.

References

[1] J. Stanberry, J.B. Balda, *A conceptual review of Sustainable Development Goal 17: Picturing politics, proximity and progress*, Journal of Tropical Futures: Sustainable Business, Governance & Development 1(1) (2024) 110-139. https://doi.org/10.1177/27538931231170509

[2] A.S. Bhatia, S. Kais, M.A. Alam, *Quantum Federated Learning in Healthcare: The Shift from Development to Deployment and from Models to Data*, IEEE Journal of Biomedical and Health Informatics 2168-2208 (2025) 1-15. https://doi.org/10.1109/JBHI.2025.3596156

[3] R. Ballester, J. Cerquides, L. Artiles, *Quantum federated learning: A comprehensive literature review of foundations, challenges, and future directions*, Quantum Machine Intelligence 7 (2025) 73. https://doi.org/10.1007/s42484-025-00292-2

[4] A. Macaluso, *Quantum supervised learning*, Künstliche Intelligenz 38 (2024) 277–291. https://doi.org/10.1007/s13218-024-00856-7

[5] S. McWeeney, T. Perciano, C. Susut, L. Chatterjee, M. Fornari, L. Biven, C. Siwy, Quantum computing for biomedical computational and data sciences: A joint DOE-NIH roundtable, U.S. Department of Energy and National Institutes of Health (2023). https://doi.org/10.2172/2228574

[6] P. Li, T. Chen, J. Liu, Enhancing quantum security over federated learning via post-quantum cryptography, arXiv (2024). https://doi.org/10.48550/arXiv.2409.04637

[7] C. Ren, Y.Z. Dong, M. Skoglund, Y. Gao, T. Wang, R. Zhang, *Enhancing dynamic security assessment in smart grids through quantum federated learning*, IEEE Transactions on Automation Science and Engineering, Advance online publication 1558-3783 (2024) 1-13. https://doi.org/10.1109/TASE.2024.3486070

[8] G. Tanbhir, M.F. Shahriyar, Quantum-inspired privacy-preserving federated learning framework for secure dementia classification, International Conference on Electrical, Computer and Communication Engineering (ECCE) (2025) 1–6. https://doi.org/10.1109/ECCE64574.2025.11013884

[9] A. Vamshikrishna, D. Ramesh, R. Mishra, N. Mohammad, Sustainable Healthcare 5.0: Integration of IoT and blockchain technology with federated learning model for securing healthcare data, Artificial intelligence of things for achieving sustainable development goals (2024) 151–170. https://doi.org/10.1007/978-3-031-53433-1_9

[10] Y. Zhou, *Quantum computing in power systems*, iEnergy 1(2) (2022) 170–187. https://doi.org/10.23919/IEN.2022.0021

[11] J.D. Fernandez, M. Brennecke, T. Barbereau, A. Rieger, G. Fridgen, Federated learning: Organizational opportunities, challenges, and adoption strategies, arXiv (2023). https://arxiv.org/abs/2308.02219

[12] S. Ashwani, A.J. Tripathy, S. Karna, P.R. ahanve, S.M. Rajagopal, Quantum computing for climate change: A comprehensive review of current applications, challenges, and future directions, 15th International Conference on Computing Communication and Networking Technologies (ICCCNT) (2024) 1–7. https://doi.org/10.1109/ICCCNT61001.2024.10724347

[13] M. Ricciardi Celsi, L. Ricciardi Celsi, *Quantum computing as a game changer on the path towards a net-zero economy: A review of the main challenges in the energy domain*, Energies 17(5) (2024) 1039. https://doi.org/10.3390/en17051039

[14] M. Ali, M. Suchismita, S.S. Ali, B.J. Choi, *Privacy-preserving machine learning for IoT-integrated smart grids: Recent advances, opportunities, and challenges*, Energies 18(10) (2025) 2515.
https://doi.org/10.3390/en18102515

[15] K. Bhattacharjya, D. De, *Federated learning-based privacy-preserving Internet of Underwater Things: A vision, architecture, computing, taxonomy, and future directions*, The Journal of Supercomputing 81 (2025) 870.
https://doi.org/10.1007/s11227-025-07322-7

[16] F. Islam, A.S. Raihan, I. Ahmed, Applications of federated learning in manufacturing: Identifying the challenges and exploring the future directions with Industry 4.0 and 5.0 visions, arXiv (2023).
https://arxiv.org/abs/2302.13514

[17] M.T.K. Ho, C.-K. Chen, L. Lee, F. Burt, S. Yu, P.H. Lee, Quantum computing for climate resilience and sustainability challenges, IEEE International Conference on Quantum Computing and Engineering (QCE) (2024)262–267.
https://doi.org/10.1109/QCE60285.2024.10289

[18] P. Priyanka Dhuliya, D.S. Rana, S. Goyal, S. Kukreti, S. Pundir, Quantum computing for sustainable development: A framework for environmental and social impact, International Conference on Advances in Computing, Communication and Materials (ICACCM) (2024) 1–7.
https://doi.org/10.1109/ICACCM61117.2024.11059008

[19] N. Arora, P. Kumar, Sustainable quantum computing: Opportunities and challenges of benchmarking carbon in the quantum computing lifecycle, arXiv (2024).
https://arxiv.org/abs/2408.05679

[20] M. Sabella, M. Vitali, Eco-friendly AI: Unleashing data power for green federated learning, arXiv (2025).
https://arxiv.org/abs/2507.17241

[21] S. Bera, T. Dey, A. Mukherjee, D. De, *IFLAG: Federated learning for sustainable irrigation in Agriculture 5.0*, IEEE Transactions on Consumer Electronics 70(1) (2024) 2303–2310.
https://doi.org/10.1109/TCE.2024.3370373