RQAI-SE: Multi-Layered Ethical and Accountable Quantum AI Engineering

Muhammad Awais, Sabeen Masood*, Farah Haneef, Taimoor Riaz, Muhammad Majid Zaman, Madiha Khanam
Dept. Software Engineering, Capital University of Science and Technology, Islamabad, Pakistan
***Correspondence**: sabeen.masood@cust.edu.pk

Citation | Awais. M, Masood. S, Haneef. F, Riaz. T, Zaman. M. M, Khanam. M, "RQAI-SE: Multi-Layered Ethical and Accountable Quantum AI Engineering", IJIST, Vol. 7 Issue. 11 pp 135-149, November 2025
Received | October 21, 2025 **Revised |** November 23, 2025 **Accepted |** November 26, 2025 **Published |** November 30, 2025.

Introduction/Importance of Study: The application of quantum Artificial Intelligence (QAI) that combines artificial intelligence and quantum computing allows simulation, optimization, and decision-making capabilities never seen before, but introduces ethical and governance issues that are beyond conventional software engineering paradigms.

Novelty statement: In this paper, we introduce the Responsible Quantum-AI Software Engineering (RQAI-SE) framework, a systematic integration of quantum-specific mechanisms for ethics, accountability, fairness, and transparency into the QAI software lifecycle. The framework comprises five hierarchically structured engineering layers, addresses four categories of quantum-specific ethical risks (algorithmic bias, lack of transparency, access inequality, and regulatory gaps), and explicitly maps its components to nine governance pillars established by the World Economic Forum (WEF) — structural contributions that distinguish RQAI-SE from existing QSE frameworks that treat ethics as secondary or post-development concerns.

Material and Method: The study employs a conceptual design methodology grounded in a systematic review of 40 high-quality sources spanning quantum software engineering, AI ethics, and governance (2021–2025). The framework is validated through comparative analysis with existing QSE frameworks Table 1 and criterion-referenced mapping to the WEF nine governance pillars Table 3.

Result and Discussion: The suggested RQAI-SE framework shows how quantum-AI systems can be designed to be both computationally efficient and responsible toward society using human-understandable quantum design, the ethical development cycle of the process, and governance models that are multidisciplinary. In contrast to the more performance-focused QSE strategies, the RQAI-SE directly deals with quantum-specific ethical risks, including but not limited to opacity, amplifying bias of probabilistic guarantees, and governance loopholes, and allows creating trustworthy, inclusive, and accountable. The model also offers a guided basis to align the emerging QAI technologies to the long-term societal values, as well as facilitating scalable and sustainable quantum software development.

Concluding Remarks: RQAI-SE provides an effective and prospective basis of responsible QAI development, under which innovations in quantum intelligence are always within the ethical, social, and governance requirements. The framework identifies five priority areas for future research, including empirical case study validation and international certification standards.

Keywords: Quantum Software Engineering; Quantum Artificial Intelligence; Responsible AI; Ethical Software Development; Governance in Quantum Systems.

Introduction:

Quantum Software Engineering (QSE) is an incipient field that generalizes software engineering to quantum computing engineering and considers the peculiarities of this type of system. One fundamental distinction between quantum and classical computation is that quantum systems can achieve exponential speedups on problems by utilizing quantum mechanical phenomena like superposition and entanglement [1]. This ability provides the need for all new approaches to system design, development, and verification that are very different than classical ones. Artificial intelligence (AI) has had many applications with QSE that have led to the development of Quantum Artificial Intelligence (QAI), which is helping to optimize machine learning models and decision-making processes with the help of quantum computing. QAI shows the transformative power in various areas such as cryptographic systems, pharmaceutical research, supply chain scientific optimization, and financial modeling [2]. In addition to acceleration of computation, qualities of QAI include the ability to search much larger solution spaces over complex optimization or simulation problems. The engineering of quantum intelligent systems is associated with significant challenges. These are the probabilistic nature of quantum systems, the instability of the hardware, the lack of good debugging tools, and the immaturity of the development infrastructure [3]. These properties render traditional deterministic software engineering unsuitable for quantum systems.

Furthermore, when used to make important judgments, QAI systems raise complex ethical and societal issues. Given the intrinsic algorithmic opacity and the concentrated global quantum computing resources, algorithmic transparency, fairness, bias minimization, and accountability are critical [4]. The existing QSE methodologies focus more on technical optimization with little concern for ethical issues [5]. Though diverse quantum development models and lifecycle paths have been discovered, the majority of them do not have a systematic way of incorporating ethical values in their formative formulations. While QSE has made significant strides in addressing the technical challenges of developing quantum applications, and the field of AI ethics has produced robust frameworks for classical systems, a critical gap exists at their intersection. Current QSE methodologies are predominantly performance-focused, treating ethical considerations such as fairness, transparency, and accountability as secondary or post-development concerns. Conversely, existing Responsible AI (RAI) frameworks are fundamentally ill-equipped to handle the novel challenges posed by quantum computing, including its probabilistic non-determinism, the opacity of quantum states, and the potential for quantum-specific bias amplification. This lack of an integrated framework creates a dangerous vacuum where powerful QAI systems could be developed and deployed without built-in ethical safeguards, leading to risks of amplified societal bias, opaque decision-making, and a lack of clear accountability. This gap is addressed in this paper through the proposed Responsible Quantum-AI Software Engineering (RQAI-SE) framework to fully integrate ethical and governance aspects of the quantum software life-cycle.

Recent advances in adjacent fields have highlighted the urgency of this gap. The integration of large language models into quantum workflows [6], the growing cybersecurity implications of quantum computing [7], and the resource-access inequalities inherent to current quantum hardware architectures [8] all point to the need for a systematic, ethics-first engineering framework for QAI development.

The RQAI-SE framework has three key contributions to the development of trustworthy QAI ecosystems: First, it designs interdisciplinary governance where the technical and ethical spaces are connected. Second, it lays down techniques of design toward ethically informed quantum systems. Third, it advances quantum technology frameworks that are accessible to solve inequalities in resource distribution. These developments counterbalance the current overemphasis on technical performance in quantum software development, which often neglects ethical considerations.

The study is guided by three explicitly defined research objectives:

To identify the quantum-specific ethical and governance gaps in existing QSE frameworks through systematic comparative analysis.

To propose an integrated, multi-layered RQAI-SE framework that embeds ethical principles, accountability mechanisms, and governance structures throughout all phases of the QAI software lifecycle.

To validate the framework's applicability by mapping its components to established global governance standards (WEF Quantum Governance Principles) and demonstrating structural differentiation from prior QSE frameworks.

Literature Review:

QSE has become a dedicated area of research that aims at modifying classical concepts of software engineering to the reality of quantum computing systems that are probabilistic and non-deterministic. In contrast to traditional software, quantum programs are subject to such fundamental principles as superposition, entanglement, and uncertainty of measurements that essentially disrupt traditional lifecycle models and verification practices [9][10].

In their initial framework research, it was noted that the classical engineering abstractions would have to be revised to provide quantum-specific properties. One of the earliest conceptual visions of QSE was introduced by [11], in which the challenges of programmability, testing, and the lifecycle management of QSE were identified. On this basis, [12] suggested a quantum software development lifecycle, reformulating requirements analysis, system design, implementation, testing, and maintenance to explicitly support quantum indeterminacy. [13] also suggested that the criteria of correctness of quantum software are fundamentally different from classical systems, which inspired new ways of validation and verification.

The verification problems had a natural consequence in architectural research in QSE. [14] suggested quantum-based architectural designs that are meant to enhance standardization, modularity, and traceability. At the level of processes, the development of hybrid quantum-classical systems has become a trend in development. Synchronized workflows to bridge the gap between classical control logic and quantum execution layers were proposed by [15], and the systematic shift from informal quantum programming practices to formal engineering approaches was reported by [16].

Quantum systems requirements engineering has also developed. [17] suggested quantum-conscious requirements models that formally include uncertainty in the functional requirements and quality requirements. Empirical studies found that there were structural obstacles to industry uptake, such as underdeveloped toolchains, lack of debugging ability, and high learning curves [18][19] These issues were further consolidated through community-driven efforts [20]. In order to overcome the practical QSE challenges, researchers have come up with various engineering solutions. Risk-aware decision-making frameworks of quantum software projects were presented by [21], and success prediction models of the project were created subsequently [22]. [23] studied the automated generation of quantum programs based on formal specifications, i.e., truth tables. Simultaneously, quantum development systems based on clouds and offering Quantum Computing as a Service (QCaaS) turned accessibility into a matter of concern but also created issues regarding reliability, security, and governance [24].

In quantum software, the recent activity shifted towards the quality assurance of its work. The methods of mutation testing [25], large-scale empirical fault studies [26], and current surveys of the state of the art in quantum software testing [27] all show that bugs and faults are still common even on an established platform like Qiskit. Specialized frameworks for the static analysis of quantum programs have also been put forward to help in the detection of defects and the reliability of the programs [28]. Decision-making and/or architectural decision-making has become a major research issue in QSE. The empirical study conducted by [29] has shown that architecture has a significant effect on reliability, maintainability, and system evolution. Nevertheless, the implicit accountability concerns are as follows, whereas even though the quality attributes of studies are discussed, the ethical implications, as well as the risks of decisions being made by AI, are not discussed as first-class architectural issues.

Simultaneously, AI has gained more and more prominence in quantum software engineering processes. Quantum circuit optimization, error mitigation, automated testing, and hybrid orchestration are now done using AI techniques [30][31]. Although such methods lead to better scalability and productivity of developers, they add extra levels of opaqueness. Based on empirical evidence, AI-assisted automation extends responsibility further between developers, AI models, and quantum execution environments, making accountability and traceability challenging [32][33]. In addition to technical aspects, the subjects of governance and ethics have become more and more popular in the environment of quantum technologies. World Economic Forum [34] suggested high-level principles to be applied to quantum computing governance and included transparency, risk management, and responsible innovation. These principles are, however, quite decoupled from tangible software engineering processes. Decision-support and risk-prediction models partially tackle the issue of governance, although not with enforcing ethical accountability; both are more focused on project outcomes than on enforced ethical accountability.

AI-driven ecosystems are also known to have researched responsible and human-centric software engineering. [35] investigated the changing role of developers within AI-enhanced settings, in which there are issues of governance and responsibility. Revisiting classical security engineering concepts in AI-dominated systems, [36] prioritized transparency, auditability, and resilience. The socio-technical aspect of quantum software development is also emphasized with diversity, inclusion, and accessibility in QSE [37][38].

More recent roadmap and vision documents unify the state of the art in QSE and explicitly demand integrated solutions that cut across the software engineering, tooling, education, ethics, and governance. Lack of cohesive frameworks was also named by [39] as a significant obstacle to reliable quantum systems. The studies mentioned above focus on service-oriented quantum computing, project governance, and the sustainability of quantum software ecosystems in the long term [40]. Although there has been significant improvement from 2022 to 2025, there is still much fragmentation in the existing studies. Ethical responsibility, accountability, and governance are not usually given serious consideration as part of engineering requirements. Specifically, the compound risks that AI-supported decision-making brings with it to quantum systems are not discussed adequately. This gap is further supported by recent demands of a quantum-oriented paradigm and ethical control.

Despite the substantial progress made in terms of quantum software lifecycles, testing, architecture, AI-assisted development, and principles of governance, there is still no single framework that deploys ethical, accountability, and governance mechanisms at all levels of Quantum AI systems. This disconnect inspires the suggested Responsible Quantum AI Software Engineering (RQAI-SE) design, which incorporates the moral, technical, and organizational ideals throughout the entire Quantum AI stack.

In spite of such technical developments, QSE research today remains dominated by engineering and infrastructure concerns, with minimal consideration of ethical issues, inclusivity, and governance mechanisms.

Table 1 demonstrates significant weaknesses across these dimensions in existing QSE frameworks, which address technical aspects thoroughly but largely ignore the socio-technical perspectives necessary for responsible quantum AI advancement. The foregoing review reveals that research has progressed along two parallel but disconnected tracks: one focused on technical rigor (lifecycles, architecture, testing) and another on high-level governance principles. There is a conspicuous absence of a cohesive engineering framework that operationalizes ethical principles such as fairness, transparency, and accountability as first-class requirements within the technical workflow of quantum software development. This paper addresses this gap by introducing the RQAI-SE framework, designed to systematically embed ethical and governance mechanisms directly into the core phases of the QAI software lifecycle, from requirements engineering to deployment and auditability.

Table 1. Comparative analysis of quantum software engineering frameworks

Framework	Focus Area	Ethical Considerations	Governance Mechanisms	Hybrid (Classical + Quantum) Support
[12]	Quantum Lifecycle	Minimal	No	Partial
[13]	Quantum Correctness	No	No	No
[14]	Quantum Architecture	No	No	Yes
RQAI-SE (Proposed)	Ethical QAI Engineering	Yes (Embedded)	Yes (Auditability, Compliance)	Yes (Full Hybrid Support)

Methodology:

The RQAI-SE framework was established to address the dual complexity of engineering and ethics in QAI systems. This multi-tiered framework enables the development of safe, transparent, and equitable QAI applications that are safe, transparent, equitable, and benefit a worldwide population by combining classical, quantum, and software engineering best practices with embedded ethical, governance, and inclusivity layers.

RQAI-SE has five interrelated layers representing each of the significant phases of the quantum software lifecycle, as shown in Figure 1. These layers span ethical requirements gathering, governance, and auditability, and are responsible for instilling responsibility into the flow of core architecture, design, and validation of the QAI systems shown in Table 2.

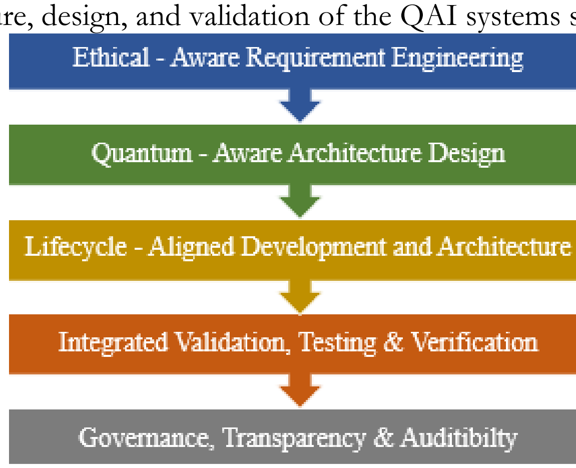

Figure 1. RQAI-SE life cycle

The Ethical-Aware Requirements Engineering layer is the first layer, which constitutes an expansion of traditional requirements engineering by integrating ethical implications in the initial phases of the design of QAI systems. As opposed to discussing fairness, bias mitigation, and accessibility as second-order issues, this layer frames them as central objectives of engineering. It promotes participatory design involving stakeholders such as legal experts, healthcare professionals, and representatives of marginalized communities, particularly in delicate areas such as criminal justice, healthcare, and money. This layer relies primarily on the processes of social and ethical risk identification, the transformation of the goal of fairness and inclusivity into technical prescriptions that can be put into practice, and a distinct definition of the extent of usage of high-hazard tasks.

Table 2. Rqai-Se Framework Layers: Goals, Activities, And Outputs

Layer	Primary Goal	Key Activities	Expected Output
Ethical-Aware Requirements	Align ethical values with system needs	Stakeholder analysis, bias identification, fairness translation	Ethics-informed requirement documents
Quantum-Aware Architecture	Ensure transparency and modular design	Risk modeling, interface traceability, and use of design patterns	Traceable hybrid architecture
Lifecycle-Aligned Development	Ethical alignment in implementation	Tool automation, low-code ethics modeling	Documented & auditable codebase
Integrated Testing & Verification	Validate fairness and accuracy	Simulation, test oracles, statistical checks	Verified quantum application
Governance & Auditability	Ensure transparency and regulatory trust	Logging, dashboards, bias audits	Compliance-ready QAI system

Quantum-Aware Architecture Design is the second layer, addressing the requirements for modular, scalable, and explainable system architectures that can host classical and quantum modules. Based on reusable architectural patterns for quantum systems, this layer underlines the need to pay attention to traceability and transparency in the interfaces and inter-component interactions of systems. Design methods entail, among other things, the use of hybrid architecture patterns, the construction of traceability matrices that map ethical demands onto particular modules, and active modelling of the way in which algorithmic biases can spread throughout the system. These measures ensure that ethical risks are considered at the beginning stages of the design process instead of reacting to them. Lifecycle-Aligned Development and Automation is the third level that aims at aligning the development practice with the ethical and functional objectives stipulated at the previous levels. This layer enables the use of model-driven engineering and platforms with low or no-code methods to facilitate faster development of QAI applications and ease traceability and reproducibility. It is implemented through ethics-aware pipelines of development, which allow checking that software artifacts comply with preset values. Also, formal documentation and automation accessories are used to making the development process reviewable and supporting post-deployment accountability.

The fourth layer, that is, Integrated Validation, Testing, and Verification, is mandatory in solving the probabilistic and non-deterministic nature of quantum computing. Traditional testing procedures cannot be adequate for QAI systems, as such systems need to have procedures that involve simulation and statistical validation. This layer adds the concept of test oracles for probabilistic systems, as well as test case generation with bias identification in mind, and composite simulation-emulation environments to test correctness and fairness during QAI decision-making. These practices are meant to enhance confidence in the performance of the systems as well as promote conformance with ethical standards. Stakeholder responsibilities by the RQAI-SE layers are represented in Figure 2. Whereas classical workflows merely trigger collaboration at opportune times, quantum workflows mandate collaborative efforts at all times- e.g., the ethicists co-design traceability in architecture (Layer 2), and the regulators validate testing (Layer 4).

The end-users also play a role at both ends (requirement and governance) to make the society aligned. The last layer, Governance, Transparency, and Auditability, has the following mechanisms of institutional accountability and regulatory compliance. Aware of the sensitivity of QAI decision-making and its societal implications, this layer inserts ethical logging, transparent dashboards, and audit trails of the bias right into the fabric of how the system operates. Along with the essential technical stakeholders, e.g., developers, ethicists, and end-users, the proposed RQAI-SE framework clearly enlarges the stakeholder model to encompass policy-makers, international standards organizations (e.g., IEEE and ISO), educators, and regulatory organizations. This expansion shows the socio-technical aspects of the Quantum Artificial Intelligence systems that exist at the edge of advanced technology, public policy, and social influences. The framework facilitates clearer role identification in other areas, including regulatory compliance, standardization, ethical oversight, and capacity building, by engaging these stakeholders throughout the various phases of the software lifecycle. Such an increased stakeholder involvement enhances accountability, promotes coordinated governance, and makes technical development consistent with legal, educational, and institutional needs throughout the lifecycle of QAI systems. Accessibility and sustainability are also considered the first-class design objectives in the RQAI-SE framework. The framework encourages equitable access to quantum technologies, and the use of cloud-based quantum computing and shared research infrastructures will provide opportunities to access quantum technologies to a wider range of people, not just the organizations that have strong quantum hardware. Concurrently, sustainability-conscious design is highlighted in light of the energy footprint of quantum hardware, the computational cost of quantum simulation, and the resource usage of large-scale hybrid QAI training protocols. Incorporating these issues into the engineering lifecycle with RQAI-SE, ethical and governance requirements are one of the main aspects of responsible QAI development that are taken into account, along with the long-term environmental and societal sustainability.

The 'multi-layered' architecture of RQAI-SE operationalizes ethical accountability through layered defense: Layer 1 prevents ethically problematic requirements from entering the system; Layer 2 ensures ethical risks are addressed in design before they are encoded in software; Layer 3 enforces ethical compliance during development; Layer 4 validates ethical compliance before deployment; and Layer 5 monitors and audits ethical compliance throughout the system's operational life. This layered approach ensures that no single point of failure can compromise the ethical integrity of the QAI system.

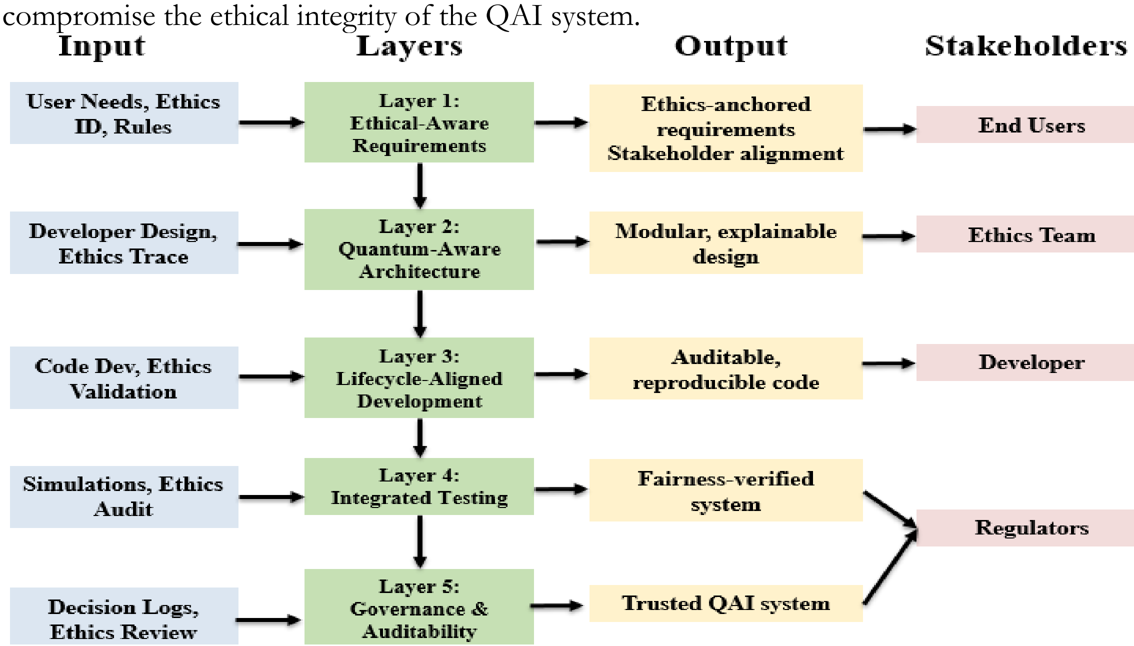

Figure 2. Stakeholder Responsibilities and Roadmap in RQAI-SE

Figure 2 illustrates the stakeholder responsibilities distributed across layers. Unlike classical workflows that trigger collaboration opportunistically, RQAI-SE mandates continuous collaboration: ethicists co-design traceability in Layer 2, regulators validate testing in Layer 4, and end-users participate in requirements (Layer 1) and governance (Layer 5).

Alignment of RQAI-SE with Global Governance Frameworks:

The layers of the RQAI-SE framework and the nine thematic governance pillars developed by the World Economic Forum (WEF).

Table 3. Mapping of RQAI-SE Framework Layers to WEF Nine Governance Pillars

RQAI-SE Layer	Primary Engineering Focus	Aligned WEF Governance Pillars	Technical–Governance Linkage
Layer 1: Ethical-Aware Requirements Engineering	Ethical risk identification, fairness goals, stakeholder inclusion	Inclusiveness & Fairness; Accountability; Human-Centric Design	Ethical objectives are translated into verifiable technical requirements, ensuring fairness, stakeholder representation, and responsible use boundaries from project inception
Layer 2: Quantum-Aware Architecture Design	Hybrid quantum–classical modular design, traceability	Transparency & Explainability; Robustness & Safety	Architectural traceability and hybrid interfaces enable explainability, fault isolation, and responsible handling of quantum uncertainty
Layer 3: Lifecycle-Aligned Development & Automation	Ethics-aware pipelines, documentation, reproducibility	Accountability; Auditability; Risk Management	Automated documentation and ethics-aligned workflows support traceable development and compliance-ready engineering practices.
Layer 4: Integrated Validation, Testing, and Verification	Probabilistic testing, bias auditing, simulation	Robustness & Safety; Reliability	Statistical testing, bias-aware test oracles, and quantum emulation enable governance-compliant validation under probabilistic execution.
Layer 5: Governance, Transparency, and Auditability	Logging, oversight, regulatory alignment	Governance & Oversight; Transparency; Societal Trust	Ethical logs, explainable dashboards, and audit trails operationalize governance requirements and enable regulatory monitoring

This mapping can be used to gain an understanding of the operationalization of the high-level governance principles, including accountability, transparency, inclusivity, and risk management, at the level of technical mechanisms throughout the quantum software lifecycle. The updated paper now shows that the ethical-conscious requirements, quantum-conscious architecture, testing, and governance layers are straightforwardly aligned to the policy goals on a global scale instead of being abstract ethical promises. As Table 3 shows, the mapping of the RQAI-SE framework shows the operationalization of the nine pillars of governance introduced by the World Economic Forum using tangible software engineering mechanisms. Instead of considering governance as an outer policy layer, RQAI-SE integrates governance principles like accountability, transparency, inclusivity, and robustness throughout every stage of the quantum-AI software lifecycle. Ethical goals are converted into technical ones, architectural traceability is used to facilitate explainability, and probabilistic testing strategies are used to guarantee quantum uncertainty robustness. This clear mapping is a way to fill a gap between principles of global governance and practical quantum software engineering practice, which is a severe weakness of current Responsible AI and Quantum Software Engineering models.

Ethical Aspects of Quantum AI Engineering:

The nature of QAI opens a new range of ethical issues that ensue. Deterministic logic, most commonly embodied in classical software engineering techniques, finds it difficult to adapt to the non-traceability and unpredictability of quantum computation. As such, these ethical aspects, as shown in Figure 3, include fairness, transparency, and accessibility as they have to be embedded in the first QAI software lifecycle stages.

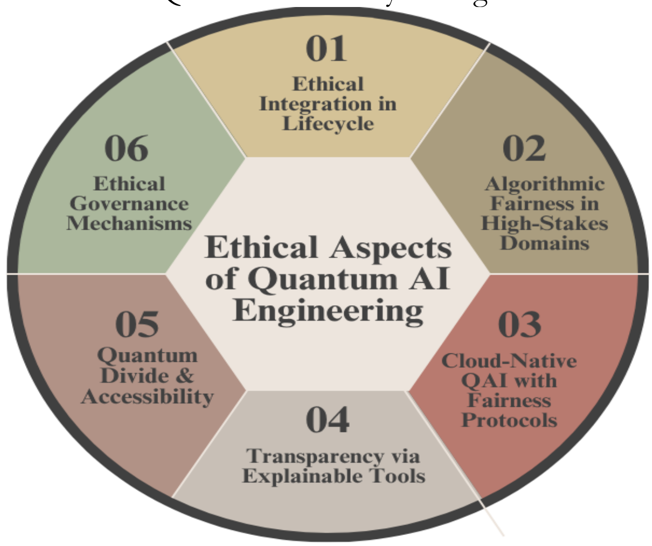

Figure 3. Ethical Aspects

One of the primary ethical concerns is algorithmic fairness, especially in fields as vital as healthcare, banking, and law. The intricacy of quantum systems' decision-making processes might exacerbate the possibility of unfair outcomes due to biased training data or design assumptions. [19] emphasize that fairness-aware design protocols are needed when integrating QIA into cloud-native environments, and these considerations should start with the requirements phase. In addition, quantum systems have a lack of transparency that cannot be resolved through regular validation and debugging strategies. According to, the tools used to achieve traceability in quantum software could be described as rare, which means that ethical logging, explainable interfaces, and simulation-based testing should be enforced to increase trust and monitoring. Table 4 presents a critical mapping of four fundamental ethical risks in quantum AI systems to their quantum-specific manifestations and the corresponding mitigation strategies within the RQAI-SE framework. This analysis reveals how quantum computing introduces novel ethical challenges that require specialized solutions.

Table 4. Ethical Risks in Quantum Ai and Rqai-Se Mitigation Strategies

Ethical Risk	Quantum-Specific Challenge	RQAI-SE Mitigation Strategy
Algorithmic Bias	Quantum noise may amplify biases	Bias-aware testing oracles (Layer 4)
Lack of Transparency	Quantum states are non-deterministic	Explainable interfaces (Layer 5)
Access Inequality	High hardware/resource barriers	Open-source tooling (Layer 1)
Regulatory Gaps	No quantum-specific policies	Ethics review boards (Layer 5)

In addition to algorithm design, QAI is concerned about fair access and control. The expertise and specially tuned hardware that quantum systems need are likely to exacerbate technological neglect in some parts of the globe. [34] prescribe open-source ecosystem, inclusive education, and participatory governance as means of achieving non-discriminatory accessibility and representation in the QAI setup. These measures will be necessary to create interdisciplinary review boards, international certification measures, and documentation requirements to help establish ethical responsibility as they grow in size.

Discussion:

The RQAI-SE framework presents a structured comparative advance over existing QSE approaches across three measurable dimensions: ethical integration depth, governance operationalization, and lifecycle coverage. Table 1 demonstrates that while existing frameworks address technical concerns competently, none systematically embeds ethical governance as an engineering requirement across all lifecycle phases. RQAI-SE addresses this structural gap through its five-layer architecture, each layer producing governance-auditable artifacts.

A structured comparison between classical AI ethics approaches and RQAI-SE reveals significant qualitative differences. Classical AI ethics frameworks primarily address data bias, model explainability, and post-deployment monitoring. Quantum AI systems require additional layers of governance: hardware-level ethical concerns arising from entanglement-induced decision opacity; bias amplification through quantum noise processes that have no classical analogue; and resource inequality challenges inherent to current quantum hardware economics. Figure 4 systematizes these differences.

Research in responsible QAI systems requires new ethical and governance models to complement classic AI solutions in their inability to cover quantum-specific issues. As systematized in Figure 4, such technological issues present themselves along three crucial axes: the complexity of technically verifying quantum probabilistic results, weaknesses in the present development pipelines regarding governance, and the inherent nature of the systemic ethical risks that occur during quantum hardware applications. Probabilistic quantum computation is the key that disturbs traditional validation paradigms. According to available testing infrastructure is inadequate to test quantum logic circuitry, and that poses a huge impediment in certifying functional correctness as well as ethical conformity. This drawback shows the urgency of finding quantum-adapted verification techniques, such as statistical test oracles verifying the probabilistic outputs, AI-based quantumized simulations of circuit behavior, and formal verification for quantum information systems. The existing lack of tough auditing mechanisms of QAI development is a problem of equal importance. A lack of thorough logging procedures and traceability attributes does make quantum systems perilous to the furtherance of hidden biases. A recent article by [32] suggests reducing these risks by using diversity-aware development toolchains, models of participatory governance, explainable quantum interfaces, and multidisciplinary ethics reviews.

Figure 4. Ethical Challenges in AI: Classical vs. Quantum Paradigms

The comparative analysis within Figure 4 allows identifying several differences between classical and quantum AI ethics. Whereas classical systems face mainly data-oriented bias and model explanations concerns, quantum systems raise the hardware-level ethical concerns, such as the decision opacity caused by entanglement, the bias amplified by quantum noise, and inequality of resources. These particular issues are taken into consideration with the RQAI-SE framework because of its quantum-specific auditing layer (Layer 4) and governance (Layer 5), which modify the classical ethical concepts to the peculiarities of quantum computing.

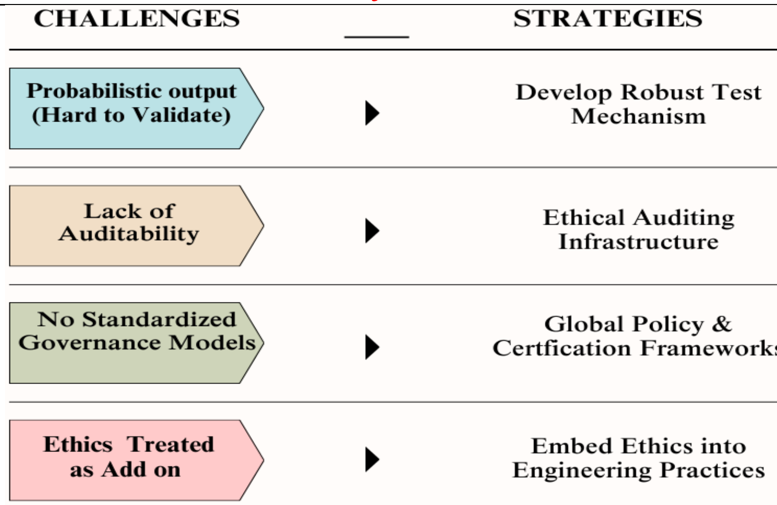

Figure 5. Challenges and Strategic Roadmap for Responsible Quantum AI Development

Figure 5 gives three strategic priorities for the responsible development of QAI. To begin with, quantum ethics should be standardized by prioritizing the formulation of quantum ethics toolkits, the generation of validation benchmark repositories, and the definition of certification procedures. Second, to implement the principles of RQAI-SE in education, it is necessary to include them in STEM curriculum, design professional education on quantum ethics, and construct ethics-engineering course content on interdisciplinary grounds. Third, international systems of governance should harmonize international standards of quantum ethics, initiate intercontinental patterns of cooperation, and globalize quantum development platforms.

To accomplish these goals, quantum engineers will have to increasingly work together with ethicists, policymakers, and people with expertise in the matter. This multidisciplinary perspective will need to balance technical advancement with wise moral judgment, and ethical concerns should be shifted past their present nature of being compliance requirements to actual design considerations. The offered roadmap focuses on proactive, not reactive, integration of ethics, making responsible innovation the center of the quantum AI development, not an addition.

The ability of quantum computers to break classical encryption poses a significant threat to global cybersecurity [7]. Despite their transformative potential, quantum computing technologies remain largely inaccessible due to high implementation costs, limited availability of hardware, and the requirement for specialized expertise [8]. To address these challenges, recent research has begun exploring the integration of Large Language Models into quantum computing workflows, particularly in areas such as quantum architecture design and quantum state simulation, to improve accessibility, automation, and efficiency in quantum system development [6]. Building upon these strategic priorities and emerging technological considerations, it is essential to acknowledge the limitations of the proposed RQAI-SE framework to guide future research and practical implementation.

The RQAI-SE framework, in its current form, is a conceptual proposition and, as such, has several limitations that must be acknowledged. First, the framework's primary limitation is its lack of empirical validation. While derived from a synthesis of existing literature, its layers, activities, and outputs have not been tested in a real-world quantum software development project. The practical feasibility, developer burden, and actual effectiveness of mechanisms like "bias-aware test oracles" or "ethics-aware development pipelines" remain unproven.

The framework’s primary limitation is the absence of empirical validation in a real-world QAI development context. As a conceptual framework derived from literature synthesis, the practical feasibility, developer burden, and effectiveness of mechanisms such as bias-aware test oracles and ethics-aware development pipelines remain to be empirically demonstrated. This limitation is characteristic of foundational framework proposals in software engineering (cf. [12][39]) and clearly identifies the primary direction for subsequent research.

Implications of the Study:

RQAI-SE contributes a new conceptual model for responsible QAI engineering that bridges the gap between technical QSE and high-level AI ethics governance. Unlike existing frameworks that treat ethics as secondary, RQAI-SE operationalizes ethical principles as first-class engineering requirements across five hierarchical layers.

For QAI development teams, the framework provides actionable layer-by-layer guidance: ethical requirements gathering (Layer 1), traceable architecture design (Layer 2), ethics-aware development pipelines (Layer 3), probabilistic testing strategies (Layer 4), and auditability mechanisms (Layer 5). Practitioners can implement these layers incrementally.

The explicit mapping of RQAI-SE layers to the WEF's nine governance pillars Table 3 offers regulators and standards organizations a concrete technical framework for policy development. This mapping demonstrates how high-level governance principles can be translated into verifiable engineering requirements.

Conclusion:

This paper introduced the RQAI-SE framework — a systematic, five-layer engineering framework for responsible Quantum AI Software Engineering. Addressing the three stated research objectives: (1) the comparative analysis in Table 1 systematically identified ethical and governance gaps across six representative QSE frameworks; (2) the five-layer RQAI-SE architecture provides an operationalized solution embedding ethics throughout the QAI lifecycle; and (3) the criterion-referenced mapping of all five layers to WEF's nine governance pillars Table 3 validates the framework's alignment with established global governance standards.

The framework's measurable structural contributions include: five hierarchically organized engineering layers; four categories of quantum-specific ethical risks with corresponding mitigation strategies; nine WEF governance pillars operationalized through concrete engineering mechanisms; and an expanded stakeholder model.

The integration of quantum computing and artificial intelligence is a revolutionary step in computational capabilities, bringing with it highly sophisticated ethical and engineering challenges that extend beyond the purview of standard software development processes. The current paper introduces the RQAI-SE framework, a systematic approach that integrates ethical frameworks such as fairness, transparency, inclusivity, and accountability in the quantum software development cycle. The framework sets the stage for the evolution of quantum AI systems that provide a balance of technological advancement with social responsibility. Future progress in responsible quantum AI development will be needed in three categories: creating specialized tools for quantum ethics verification and bias identification, establishing global certification requirements for quantum AI systems, and creating regulatory frameworks addressing the distinct nature of quantum computing. Effective progress of quantum AI technologies relies on interdisciplinary collaboration between quantum engineering, ethics, policymaking, and domain-specific applications. The RQAI-SE framework provides the methodological foundation to achieve that necessary balance between technological innovation and ethical responsibility in quantum computing applications.

Five structured priorities for future research are identified:

Empirical validation of RQAI-SE through real-world QAI development case studies across multiple domains (healthcare, finance, criminal justice), measuring developer burden, compliance effectiveness, and ethical risk mitigation outcomes.

Development of quantum ethics toolkits, including bias-aware test oracle libraries, ethical logging specifications, and explainability dashboard templates compatible with leading quantum computing platforms (Qiskit, Cirq, Azure Quantum).

Establishment of international certification standards for Quantum AI systems aligned with RQAI-SE's governance layer requirements and WEF governance pillars.

Integration of RQAI-SE principles into STEM curricula and professional development programs.

Development of cross-jurisdictional quantum governance frameworks harmonizing divergent national regulatory approaches to QAI accountability and access equity.

Acknowledgement: The authors would like to recognize those in the academic community and previous researchers whose scientific contributions in the field of quantum computing, artificial intelligence, and software engineering led to the creation of the present study. The constructive ideas of the fellow scholars are also valued by the authors, and they assisted in polishing the conceptual framework and presentation of this research. The authors declare that this manuscript has not been previously published and is not under consideration for publication in any other journal.

Author's Contribution: Every author participated equally in this study. All authors participated in the study conceptualization, literature review, design of the structure, preparation of the manuscript, and critical review of the content. The submission process and coordination of the research process were handled by the corresponding author. The final manuscript was reviewed and approved by all the authors.

Conflict of Interest: The authors declare that they do not have a conflict of interest when publishing this manuscript in the International Journal of Information Science and Technology (IJIST).

Project Details: This study was not an undertaking in any funded or registered research. Thus, it does not have a project number, project cost, or completion date.

References:

[1] Aritra Sarkar, "Automated quantum software engineering," *Autom. Softw. Eng.*, vol. 31, no. 36, 2024, [Online]. Available: https://link.springer.com/article/10.1007/s10515-024-00436-x

[2] G. Bisicchia, J. García-Alonso, J. M. Murillo, and A. Brogi, "From Quantum Software Handcrafting to Quantum Software Engineering," *Proc. - 2024 IEEE Int. Conf. Softw. Anal. Evol. Reengineering - Companion, SANER-C 2024*, pp. 149–150, 2024, doi: 10.1109/SANER-C62648.2024.00026.

[3] Alvaro M. Aparicio-Morales, Enrique Moguel, Luis Mariano Bibbo, Alejandro Fernandez, Jose Garcia-Alonso, Juan M. Murillo, "An Overview of Quantum Software Engineering in Latin America," *arXiv:2405.20661*, 2024, [Online]. Available: https://arxiv.org/abs/2405.20661

[4] Francisco Valdes-Souto, Hector G. Perez-Gonzalez, Carlos A. Perez-Delgado, "Q-COSMIC: Quantum Software Metrics Based on COSMIC (ISO/IEC19761)," *arXiv:2402.08505*, 2024, [Online]. Available: https://arxiv.org/abs/2402.08505

[5] Kanishk Dwivedi, Majid Haghparast & Tommi Mikkonen, "Quantum software engineering and quantum software development lifecycle: a survey," *Cluster Comput.*, vol. 27, pp. 7127–7145, 2024, [Online]. Available: https://link.springer.com/article/10.1007/s10586-024-04362-1

[6] Shuangxiang Zhou, Ronghang Chen, Zheng An, Shi-Yao Hou, "Application of Large Language Models to Quantum State Simulation," *arXiv:2410.06629*, 2024, [Online]. Available: https://arxiv.org/abs/2410.06629

[7] S. Paul, N. R. Choudhury, B. Pandit, and A. Dawn, "Integration of AI and Quantum Computing in Cybersecurity: A Comprehensive Review," *Integr. AI, Quantum Comput. Semicond. Technol.*, pp. 287–308, Jan. 2024, doi: 10.4018/979-8-3693-7076-6.ch014.

[8] D. Chauhan, P. Ranka, P. Bahad, and R. Pathak, "Applications of Quantum Artificial Intelligence: A Systematic Review," *Integr. AI, Quantum Comput. Semicond. Technol.*, pp. 159–182, Jan. 2024, doi: 10.4018/979-8-3693-7076-6.ch008.

[9] T. Y. Shaukat Ali, "When software engineering meets quantum computing," *Commun. ACM*, vol. 65, no. 4, pp. 84–88, 2022, [Online]. Available: https://dl.acm.org/doi/10.1145/3512340

[10] Diego Alonso, Pedro Sánchez, "Engineering the development of quantum programs: Application to the Boolean satisfiability problem," *Adv. Eng. Softw.*, vol. 173, p. 103216, 2022, doi: https://doi.org/10.1016/j.advengsoft.2022.103216.

[11] M. Piattini, M. Serrano, R. Perez-Castillo, G. Petersen, and J. L. Hevia, "Toward a Quantum Software Engineering," *IT Prof.*, vol. 23, no. 1, pp. 62–66, Jan. 2021, doi: 10.1109/MITP.2020.3019522.

[12] Benjamin Weder, Johanna Barzen, Frank Leymann, Daniel Vietz, "Quantum Software Development Lifecycle," *arXiv:2106.09323*, 2021, [Online]. Available: https://arxiv.org/abs/2106.09323

[13] J. Zhao, "Quantum Software Engineering: Landscapes and Horizons," *arXiv:2007.07047*, 2020, [Online]. Available: https://arxiv.org/abs/2007.07047

[14] M. Baczyk, R. Pérez-Castillo, and M. Piattini, "Towards a Framework of Architectural Patterns for Quantum Software Engineering," *Proc. - IEEE Quantum Week 2024, QCE 2024*, vol. 2, pp. 228–233, 2024, doi: 10.1109/QCE60285.2024.10283.

[15] S. Nagarajan, M. Malarvel, and J. Thangakumar, "A Framework for Quantum based Software Development Process," *2024 Int. Conf. Adv. Data Eng. Intell. Comput. Syst. ADICS 2024*, 2024, doi: 10.1109/ADICS58448.2024.10533595.

[16] A. B. Giuseppe Bisicchia, Jose Garcia-Alonso, Juan M. Murillo, "From Quantum Mechanics to Quantum Software Engineering: A Historical Review," *arXiv:2404.19428*, 2024, [Online]. Available: https://arxiv.org/abs/2404.19428

[17] P. Spoletini, "Towards Quantum Requirements Engineering," *Proc. - 31st IEEE Int. Requir. Eng. Conf. Work. REW 2023*, pp. 371–374, 2023, doi: 10.1109/REW57809.2023.00072.

[18] M. R. El Aoun, H. Li, F. Khomh, and M. Openja, "Understanding Quantum Software Engineering Challenges An Empirical Study on Stack Exchange Forums and GitHub Issues," *Proc. - 2021 IEEE Int. Conf. Softw. Maint. Evol. ICSME 2021*, pp. 343–354, 2021, doi: 10.1109/ICSME52107.2021.00037.

[19] Moses Openja, Mohammad Mehdi Morovati, "Technical debts and faults in open-source quantum software systems: An empirical study," *J. Syst. Softw.*, vol. 193, p. 111458, 2022, doi: https://doi.org/10.1016/j.jss.2022.111458.

[20] Michael Felderer, Davide Taibi, "Software Engineering Challenges for Quantum Computing: Report from the First Working Seminar on Quantum Software Engineering (WSQSE 22)," *ACM SIGSOFT Softw. Eng. Notes*, vol. 48, no. 2, pp. 29–32, 2023, [Online]. Available: https://dl.acm.org/doi/abs/10.1145/3587062.3587071

[21] A. A. K. Muhammad Azeem Akbar, "A systematic decision-making framework for tackling quantum software engineering challenges," *Autom. Softw. Eng.*, vol. 30, 2023, [Online]. Available: https://link.springer.com/article/10.1007/s10515-023-00389-7

[22] Muhammad Azeem Akbar, Arif Ali Khan, "Genetic model-based success probability prediction of quantum software development projects," *Inf. Softw. Technol.*, vol. 165, p. 107352, 2024, doi: https://doi.org/10.1016/j.infsof.2023.107352.

[23] Macario Polo, Ignacio García-Rodríguez, Manuel Serrano, Mario Piattini, "Generation of Quantum Software From Truth Tables," *Softw. Pract. Exp.*, vol. 55, no. 8, 2025, [Online]. Available: https://onlinelibrary.wiley.com/doi/10.1002/spe.3426?af=R

[24] Aakash Ahmad, Muhammad Waseem, Peng Liang, Mahdi Fehmideh, Arif Ali Khan, David Georg Reichelt, Tommi Mikkonen, "Engineering Software Systems for Quantum Computing as a Service: A Mapping Study," *arXiv:2303.14713*, 2023, [Online]. Available:

https://arxiv.org/abs/2303.14713
[25] D. Fortunato, J. CAMPOS and R. ABREU, "Mutation Testing of Quantum Programs: A Case Study With Qiskit," *IEEE Trans. Quantum Eng.*, vol. 3, pp. 1–17, 2022, doi: 10.1109/TQE.2022.3195061.
[26] Matteo Paltenghi, Michael Pradel, "Bugs in Quantum computing platforms: an empirical study," *Proc. ACM Program. Lang.*, 2022, [Online]. Available: https://dl.acm.org/doi/10.1145/3527330
[27] M. P. Antonio García de la Barrera, Ignacio García-Rodríguez de Guzmán, Macario Polo, "Quantum software testing: State of the art," *J. Softw. Evol. Process*, 2023, [Online]. Available: https://onlinelibrary.wiley.com/doi/10.1002/smr.2419
[28] Matteo Paltenghi, Michael Pradel, "Analyzing Quantum Programs with LintQ: A Static Analysis Framework for Qiskit," *Proc. ACM Softw. Eng.*, 2024, [Online]. Available: https://dl.acm.org/doi/10.1145/3660802
[29] Mst Shamima Aktar, Peng Liang, "Architecture decisions in quantum software systems: An empirical study on Stack Exchange and GitHub," *Inf. Softw. Technol.*, vol. 177, p. 107587, 2025, doi: https://doi.org/10.1016/j.infsof.2024.107587.
[30] Shaukat Ali, Tao Yue, "On the Need of Quantum-Oriented Paradigm," *QP4SE 2023 - Proc. 2nd Int. Work. Quantum Program. Softw. Eng. Co-located with ESEC/FSE 2023*, 2023, [Online]. Available: https://dl.acm.org/doi/10.1145/3617570.3617868
[31] Asmar Muqeet, Tao Yue, "Quantum Software Testing 101," *Proc. - Int. Conf. Softw. Eng.*, 2024, [Online]. Available: https://dl.acm.org/doi/10.1145/3639478.3643059
[32] Muneera Bano, Shaukat Ali, Didar Zowghi, "Envisioning Responsible Quantum Software Engineering and Quantum Artificial Intelligence," *arXiv:2410.23972*, 2025, [Online]. Available: https://arxiv.org/abs/2410.23972
[33] Valerio Terragni, Partha Roop, Kelly Blincoe, "The Future of Software Engineering in an AI-Driven World," *arXiv:2406.07737*, 2024, [Online]. Available: https://arxiv.org/abs/2406.07737
[34] "Rigene Project - Quantum computing Principles." Accessed: Apr. 22, 2026. [Online]. Available: https://www.rigeneproject.org/cognitive-news/quantum-computing-principles
[35] Ketai Qiu, Niccolò Puccinelli, "From Today's Code to Tomorrow's Symphony: The AI Transformation of Developer's Routine by 2030," *ACM Trans. Softw. Eng. Methodol.*, vol. 34, no. 5, pp. 1–17, 2025, [Online]. Available: https://dl.acm.org/doi/10.1145/3709353
[36] Nikhil Patnaik, Joseph Hallett, Awais Rashid, "Saltzer & Schroeder for 2030: Security engineering principles in a world of AI," *arXiv:2407.05710*, 2024, [Online]. Available: https://arxiv.org/abs/2407.05710
[37] Enrique Moguel, Javier Rojo, David Valencia, Javier Berrocal, Jose Garcia-Alonso & Juan M. Murillo, "Quantum service-oriented computing: current landscape and challenges," *Softw. Qual. J.*, vol. 30, pp. 983–1002, 2022, [Online]. Available: https://link.springer.com/article/10.1007/s11219-022-09589-y
[38] Muhammad Azeem Akbar, Arif Ali Khan, Sajjad Mahmood, Saima Rafi, "Quantum Software Engineering: A New Genre of Computing," *arXiv:2211.13990*, 2022, [Online]. Available: https://arxiv.org/abs/2211.13990
[39] J. G.-A. Juan Manuel Murillo, "Quantum Software Engineering: Roadmap and Challenges Ahead," *ACM Trans. Softw. Eng. Methodol.*, vol. 34, no. 5, 2025, [Online]. Available: https://dl.acm.org/doi/10.1145/3712002
[40] N. Arora and P. Kumar, "Sustainable Quantum Computing," *Commun. ACM*, vol. 69, no. 1, pp. 84–90, Jan. 2026, doi: 10.1145/3745782.