Brief Review: Artificial intelligence in the age of quantum computing

JACOB NAGLER[1,*], OFIR BEN-SHABAT[2]
[1]Researcher Engineer
NIRC,
Haifa, Givat Dowens
ISRAEL

[2]Senior DevOps Engineer,
Amdocs, Ra`anana,
ISRAEL

[*]*Corresponding Author*

Abstract- This paper presents a rigorous, evidence-driven synthesis of the technical and socio-technical landscape at the intersection of artificial intelligence (AI) and quantum computing (QC). Combining a structured literature synthesis with a transparent diagnostic scoring framework, the study evaluates ten principal subdomains, including quantum hardware and control, quantum error correction, hybrid classical-quantum architectures, quantum machine learning (QML), quantum neural networks (QNNs), optimization methods, data-encoding strategies, post-quantum cryptography, software/benchmarking, and ethics/governance along three axes: readiness, technical risk, and potential impact. Key findings identify quantum hardware and error correction as foundational enablers with the greatest transformative potential but minimal near-term readiness; by contrast, hybrid architectures and quantum-safe cryptographic measures exhibit higher maturity and represent the most practicable near-term routes to benefit from NISQ devices. QML and QNN approaches retain considerable theoretical promise but are constrained by data-encoding costs, training-landscape pathologies (e.g., barren plateaus), noise sensitivity, and dequantizing caveats. To move from theoretical potential to verifiable progress, the paper proposes a prioritized, technical agenda that emphasizes (i) standardized, reproducible benchmarking pipelines and open data; (ii) hybrid co-design experiments to disentangle expressive from noise; (iii) resource-aware error-mitigation and low-overhead QEC research; and (iv) explicit articulation of input/state-preparation models in all speedup claims. The combined technical survey, diagnostic dataset, and concrete experimental protocols aim to enable empirically testable claims of quantum-augmented AI and to guide coordinated research and engineering efforts toward fault-tolerant, socially responsible outcomes.

Key-words: Quantum computing; Quantum machine learning; Variational algorithms; Hybrid architectures; Quantum error correction; Post-quantum cryptography; Benchmarking; NISQ.

Received: May 29, 2025. Revised: September 15, 2025. Accepted: October 16, 2025. Published: April 7, 2026.

1. Introduction

Artificial intelligence (AI) has become ubiquitous, driving advances in sectors such as healthcare, finance, manufacturing, transportation and social media. AI systems are built using algorithms that learn from data, adapt to changing environments and make decisions or predictions. However, many AI tasks; e.g., such as training deep neural networks [1], solving combinatorial optimization problems or processing natural language are computationally intensive and push the limits of classical computing resources. Quantum computing [2], an emerging paradigm based on principles of quantum mechanics, offers the possibility of overcoming these limitations. Whereas classical computers encode information in bits that take values 0 or 1, quantum computers use qubits that can exist in super-positions of states and can become entangled [3]. These properties enable quantum computers to explore exponentially large state spaces with fewer resources than classical computers [4]. Quantum algorithms such as Grover's search [5] provide quadratic speed ups for unstructured search problems [4], and Shor's factorization algorithm [6] achieves polynomial time factorization, threatening widely used cryptographic systems [7]-[8].

The age of quantum computing [2] is arriving. Google's 53 qubit Sycamore processor demonstrated a quantum computation that required only minutes but would take days on the most powerful classical supercomputer [9]. IBM has announced 127 qubit and 433 qubit devices, with roadmaps targeting thousands of qubits. Meanwhile, neutral atom and photonics platforms are making progress toward error corrected quantum machines. Yet current devices belong to the noisy intermediate scale quantum (NISQ) era, characterized by limited qubit counts, short coherence times and imperfect gates [4]. Exploiting NISQ devices for practical advantage requires hybrid quantum–classical algorithms [10] and error mitigation techniques. This context motivates the integration of AI and quantum computing.

AI in the age of quantum computing manifests in two directions: (1) using quantum computers to accelerate or enhance AI algorithms (quantum for AI), and (2) applying AI techniques to design, control and improve quantum computers and algorithms (AI for quantum). The synergy between these directions may catalyze advances in both fields. Quantum machine learning promises to speed up data processing and pattern recognition by leveraging quantum state space [9]-[11]. Quantum reinforcement learning aims to learn optimal policies for decision making under uncertainty, with potential speed ups and novel behaviors. Conversely, AI helps calibrate and optimize quantum hardware, identify error models and even discover new quantum algorithms ([7], [12]). Policy analysts recognize that the convergence of quantum technology and AI offers opportunities but raises risks, including new cybersecurity threats, inequalities in access and ethical dilemmas [13]. Understanding this emerging landscape requires a holistic analysis across technical, ethical and societal dimensions.

Our approach synthesizes key technical results on NISQ-era algorithms and their limitations (e.g., the NISQ concept and cautious expectations set out by Preskill), and integrates contemporary analyses of variational algorithms, dequantizing results and training-landscape pathology (barren plateaus) ([14]; [15], [16]).

This paper provides a comprehensive survey and analysis of AI in the age of quantum computing. Section 2 reviews quantum computing fundamentals and complexity theory relevant to AI. Section 3 covers quantum machine learning, including variational circuits, quantum kernel methods and generative models. Section 4 discusses quantum reinforcement learning and quantum agents. Section 5 explores quantum natural language processing and quantum computer vision. Section 6 examines AI for quantum applications such as hardware design, control and error correction. Section 7 analyses quantum optimization, including quantum annealing and the quantum approximate optimization algorithm (QAOA). Section 8 discusses algorithmic complexity and fairness, highlighting dequantizing results and quantum enhanced fairness frameworks. Section 9 investigates ethical, policy and sustainability issues. Section 10 synthesizes challenges and opportunities, suggesting future research directions.

The diagram exhibited in Fig. 1 schematically depicts the intersection between quantum computing (left) and artificial intelligence (right) and highlights four principal translational pathways examined in the manuscript: quantum machine learning (novel feature maps and kernel methods), classical–quantum hybrid architectures (variational and orchestration frameworks), quantum error correction and mitigation (resource-aware fault-tolerance strategies), and post-quantum cryptography (standards and migration). Arrows and data-flow motifs indicate bidirectional coupling classical pre/post-processing integrated with quantum subroutines, while the layout emphasizes (i) the centrality of hybrid co-design for near-term experiments, (ii) the foundational role of hardware and error control for long-term impact, and (iii) the operational and security considerations required for responsible deployment.

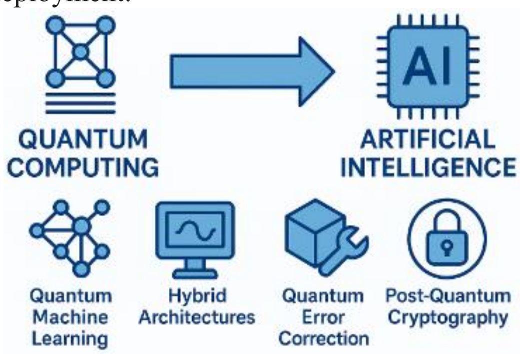

Fig. 1: Conceptual overview of the manuscript's scope and research themes.

2. Quantum Computing Fundamentals and Complexity Theory

2.1 Qubits, Superposition and Entanglement

Quantum computation [2] operates on qubits, which are two level quantum systems whose state can be written as $|\psi\rangle = \alpha|0\rangle + \beta|1\rangle$, where α and β are complex amplitudes satisfying $|\alpha|^2 + |\beta|^2 = 1$. The principle of *superposition* allows qubits to represent combinations of states [3], enabling the exploration of many computational paths simultaneously [4]. Measurement collapses the superposition to a classical outcome with probabilities proportional to the squared amplitudes. Another central resource is *entanglement*, where the joint state of multiple qubits cannot be factorized into independent single qubit states. Entangled qubits exhibit correlations that persist regardless of spatial separation, providing advantages in algorithms, communication and cryptography. Quantum gates are reversible unitary operations acting on qubits; sequences of gates form quantum circuits [17]. The *Hadamard gate* creates superposition, the CNOT gate entangles qubits, and phase gates adjust amplitudes. Quantum circuits can implement complex algorithms by exploiting interference between computational paths. However, qubits are susceptible to decoherence and noise, requiring error mitigation or error correction for reliable computation.

2.2 Quantum Algorithms: Shor and Grover

The potential of quantum computing was first illustrated by algorithms that outperform classical ones. Shor's algorithm [6] for integer factorization uses the quantum Fourier transform to find periods of modular exponentiation, yielding polynomial-time complexity, whereas classical algorithms [10] are sub-exponential [8]. Breaking RSA encryption with Shor's algorithm would compromise much of today's digital security. Grover's search algorithm provides a quadratic speed-up for unstructured search. Given an unsorted database of N items, Grover's algorithm finds a marked item in $O(\sqrt{N})$ steps by alternating an oracle that flips the phase of the marked state with a diffusion operation that amplifies its amplitude [4]. Although the speed-up is quadratic rather than exponential, Grover's algorithm is widely applicable and underpins quantum search subroutines used in optimization and machine learning [5].

These algorithms motivate the exploration of *quantum complexity classes*. The class BQP (bounded-error quantum polynomial time) comprises decision problems solvable by a quantum computer in polynomial time with bounded error. It is known that $P \subseteq BPP \subseteq BQP \subseteq PP \subseteq PSPACE$ [18]. Factorization is in BQP, but whether BQP extends beyond NP or collapses to BPP remains open. Oracle results indicate that BQP is not contained in BPP or in the polynomial hierarchy relative to certain oracles [18]. Quantum algorithms such as *quantum annealing* [19], *quantum walks* and *amplitude amplification* provide various speed-ups, but not every problem benefits. Evaluating purported quantum advantages requires careful analysis of input assumptions and complexity theory.

2.3 Quantum Computing Hardware and the NISQ Era

Practical quantum computing faces hardware challenges. Early demonstrations such as IBM's *Q System One* and Google's *Sycamore* achieved around 50–60 qubits [9], while neutral atom platforms, trapped ions and photonic systems promise higher fidelities. Qubits are implemented using superconducting circuits ([20]-[21]), trapped ions, semiconductor spins, photons or neutral atoms. Each platform offers trade-offs in coherence time, gate speed and connectivity. Current devices suffer from decoherence, gate errors and limited qubit connectivity; thus, algorithms must be designed to tolerate noise. The NISQ era emphasizes hybrid quantum–classical algorithms [10] that run shallow quantum circuits embedded in classical workflows. These include variational quantum Eigensolvers (VQEs) [22], quantum approximate optimization algorithms (QAOA) and quantum machine-learning models. Advances in quantum hardware, error correction and control will determine when true fault-tolerant quantum computing arrives.

3 Quantum Machine Learning

Quantum machine learning (QML) merges quantum computing with artificial intelligence in order to solve computationally intensive learning tasks more efficiently [23]. Quantum artificial intelligence is a broader term that includes quantum reasoning, planning, natural language processing and agents ([11], [24]). QML aims to harness quantum superposition and entanglement to encode and process data in high-dimensional Hilbert spaces [25], potentially achieving speed-ups over classical algorithms [9]. However, claims of exponential speed-ups require careful scrutiny, as many depend on strong assumptions about efficient state preparation or access models (e.g., QRAM), which can also empower classical sampling methods [26]-[27]. In this section, QML models, including quantum kernel methods, variational circuits, quantum generative models and quantum approximate optimization are reviewed.

3.1 Quantum Kernel Methods and Support Vector Machines

Classical quantum support vector machines (QSVMs) [28] perform classification compared to the classical SVMs [29] by computing inner products between data points in a high-dimensional feature space using a kernel function. *Quantum kernel methods* leverage quantum circuits to embed classical data into quantum states and compute inner products (kernel values) via measurements. A *quantum support vector machine* (QSVM) uses quantum feature maps to classify data by maximizing the margin between classes. The QSVM defines a quantum kernel $\kappa(x, y) = |\langle\Phi(x)|\Phi(y)\rangle|^2$, where $|\Phi(x)\rangle$ is a quantum state encoding the data point x. In an experiment on trapped-ion hardware, a quantum SVM performed classification and regression tasks with accuracy comparable to noiseless simulations [30]. Low-rank approximations of quantum kernels improved regression performance on NISQ devices [30]. Quantum kernels may offer advantages for problems that are hard to approximate classically; however, efficient data encoding and circuit depth remain challenges.

3.2 Variational Quantum Circuits and Quantum Neural Networks

Variational quantum circuits (VQCs) are parameterized circuits trained using classical optimization. A cost function such as classification loss or energy expectation (World Economic Forum 2025 [31]) is evaluated on a quantum computer, and a classical optimizer updates the parameters to minimize it. VQCs can approximate functions and implement quantum neural networks [1]. They underpin algorithms such as the variational quantum Eigensolver (VQE) for chemistry ([16], [32]) and the *quantum approximate optimization algorithm* (QAOA) for combinatorial problems. VQCs are suitable for NISQ hardware because they can be shallow and tolerate noise, but they suffer from *barren plateaus* [15]. Formally, for a wide class of parameterized quantum circuits, the variance of the partial derivative of the cost function C with respect to a parameter θ_k decays exponentially with the number of qubits n : $\mathrm{Var}[\partial_{\theta_k} C] \in \mathcal{O}(1/b^n)$ for some base $b > 1$. This exponential decay of the gradient landscape implies that an exponentially large number of measurements is required to distinguish the update direction from statistical noise, thereby hindering training. Techniques like layer-wise learning, parameter initialization and tailored ansatz seek to mitigate barren plateaus.

Techniques like layer wise learning, parameter initialization and tailored ansatz seek to mitigate barren plateaus. Hybrid architectures that combine classical neural networks with VQCs can reduce the number of quantum parameters [1].

3.3 Quantum Generative Models

Generative models learn to sample from complex probability distributions. Quantum systems naturally represent probability distributions via measurement outcomes. *Quantum generative adversarial networks* (QGANs) combine a quantum generator with a quantum or classical discriminator. A survey of QGAN architectures distinguishes fully quantum QGANs, where both generator and discriminator are quantum, from hybrid QGANs with one classical component. Applications include image generation, biology problems [33] drug discovery and finance [34]. The survey notes that QGANs can approximate distributions more efficiently than classical GANs for some tasks, but they face challenges such as noise, barren plateaus and limited qubit numbers [34]. Another quantum generative model is the *quantum circuit Born machine* (QCBM), which samples from distributions encoded in a quantum circuit's Born rule. QCBMs have been used to learn classical data distributions and may provide advantages in modeling high-dimensional data.

3.4 Quantum Principal Component Analysis and Dequantizing

Quantum principal component analysis (QPCA) seeks to find eigenvectors of a density matrix using phase estimation, promising exponential speed-ups. However, recent work by Tang [26] shows that when the classical algorithm is given access to sample and query (SQ) access to data - analogous to the state preparation assumptions of QPCA, the exponential speed-ups vanish [26]. Tang dequantizes QPCA by demonstrating that classical algorithms can achieve recommendation and clustering tasks in polynomial time. Specifically, Tang's classical analogue runs in $\mathcal{O}(\frac{k^6\|A\|_F^6}{\epsilon^6})$ operations for a rank-k matrix A and error margin ϵ, directly replacing the previously assumed exponential quantum run-time of $\mathcal{O}(\text{poly}(k)\log d)$ with a classical polynomial factor. This result underscores the importance of carefully comparing quantum algorithms to appropriately empowered classical algorithms [7]. Many quantum linear algebra algorithms rely on efficient state preparation, which may be as hard as the original problem.

3.5 Quantum Approximate Optimization Algorithm

The *quantum approximate optimization algorithm* (QAOA) is a variational algorithm for solving combinatorial optimization problems. QAOA alternates between applying a cost Hamiltonian that encodes the problem and a mixing Hamiltonian that introduces transitions between states. The depth *p* controls the number of alternations; as p increases, QAOA approaches the true ground state. In a study of QAOA applied to the Sherrington–Kirkpatrick spin glass model, the energy achieved at depth $p = 11$ was shown to exceed the best classical semidefinite programming algorithm [35]. Importantly, the authors derived a formula that allows the energy expectation of large systems to be computed analytically, enabling parameter optimization without simulation [35]. QAOA is widely studied for Max-Cut, traveling salesperson, scheduling and portfolio optimization. However, its performance on real-world instances is still being characterized, and training parameters can be difficult for high p due to barren plateaus.

3.6 Quantum Annealing and Optimization

Quantum annealing uses quantum fluctuations to find the ground state of a problem Hamiltonian [19]. It starts from an initial Hamiltonian with an easily prepared ground state and slowly interpolates to the problem Hamiltonian, relying on the adiabatic theorem. *Quantum annealers* solve quadratic unconstrained binary optimization (QUBO) problems by mapping them to Ising models. Early devices showed limited advantage, but a recent benchmarking study using a >5,000 qubit quantum annealer with improved connectivity (Pegasus topology) found that the annealer solved large, dense QUBO problems

6,561 times faster than classical solvers and produced higher-quality solutions [36]. Hybrid approaches such as Qbsolv decompose large QUBO instances into sub-problems solved on the annealer [36]. Quantum annealing is applied to vehicle routing, portfolio optimization, machine learning and scheduling [19]. However, its true complexity advantage remains under study, and noise, limited connectivity and embedding overhead are challenges.

4 Quantum Reinforcement Learning

Reinforcement learning (RL) studies how agents learn policies through interactions with an environment to maximize cumulative reward. Quantum reinforcement learning (QRL) explores RL in quantum environments, RL using quantum information processing or quantum enhancements to classical RL [37]. A survey of 177 QRL papers classifies approaches by degree of quantization: *quantum-inspired RL* uses quantum probabilistic principles; *variational quantum circuit (VQC) RL* implements policy or value functions using parameterized quantum circuits; *post-NISQ QRL* includes algorithms requiring fault-tolerant quantum computation; and *quantum RL applications* apply RL to quantum systems [38]. The survey notes that 22 works are quantum-inspired, 68 use VQC, 26 are applications and 30 are post-NISQ; it emphasizes that guaranteed quantum advantage is proven only for post-NISQ algorithms [38]. VQC-based QRL leverages quantum circuits to represent policies or value functions; training involves measuring expectation values and updating parameters via classical optimizers. Quantum states may encode superposition of actions, enabling exploration over multiple actions at once. However, QRL suffers from barren plateaus, data loading issues and the need for many samples.

In addition to quantum-for-RL, RL can enhance quantum computing. For example, a reinforcement learning framework optimizes quantum error-correcting codes by modifying surface codes to minimize logical error rates [39]. The RL agent explores modifications to the code geometry and learns to find near-optimal layouts. This approach reduces the resources required for error correction and transfers knowledge across different noise models, demonstrating AI's power to improve quantum hardware reliability. RL has also been used to optimize pulse sequences and control parameters in quantum devices [12].

Quantum RL also includes models where the environment or agent is quantum. For instance, *quantum projective simulation* uses quantum random walks in the agent's memory to improve decision making. *Quantum tabular RL* uses quantum probability amplitudes to speed up value-function updates. Some approaches propose using quantum annealers to accelerate RL by encoding the policy search as QUBO. Nevertheless, no general quantum advantage has been demonstrated for RL in the NISQ era, and more research is needed.

5 Quantum Natural Language Processing and Computer Vision

5.1 Quantum Natural Language Processing

Natural language processing (NLP) analyses and generates human language. *Quantum natural language processing* (QNLP) integrates quantum computing with NLP, aiming to process complex linguistic structures more efficiently [40]. A 2024 survey notes that QNLP is one of several subfields of quantum AI alongside quantum reasoning, automated planning, computer vision and quantum agents [24]. An introductory article on QNLP for bioinformatics explains that QNLP uses quantum principles such as superposition and entanglement to represent linguistic structures, enabling quantum systems to handle the combinatorial complexity of biological text ([33], [41]). Applications include genomic sequence analysis, protein structure prediction and drug discovery ([41]). QNLP methods may improve tasks like protein folding prediction and ligand binding constant estimation by exploiting quantum parallelism. However, challenges include encoding classical text into quantum states, limited qubit numbers and noisy hardware ([41]).

Near-term experiments show modest progress. A study performed the largest quantum NLP experiment to date by encoding representations of approximately 10,000 words into quantum states and performing bigram classification on a quantum computer, achieving about 62 % accuracy [40]. They used a quantum circuit Born machine and quantum probability for bigram modeling. The results highlight that quantum NLP experiments currently handle small datasets and exhibit high variance across real datasets [40]. Future work may involve hybrid quantum–classical models for semantic analysis, syntax parsing and machine translation. The interplay between quantum semantics (e.g., distributional compositional semantics) and quantum hardware remains an exciting research frontier.

5.2 Quantum Computer Vision

Computer vision enables machines to interpret visual information [37]. Quantum computer vision explores using quantum computation to accelerate image processing, feature extraction and classification. Quantum algorithms for Fourier transforms, wavelet transforms and edge detection can process images in superposition. Quantum image representation frameworks such as *Flexible Representation of Quantum Images* (FRQI) encode pixel intensities and positions into amplitudes and phases. Quantum convolutional neural networks (QCNNs) adapt classical convolution layers to quantum circuits. Quantum generative models can synthesize images with fewer parameters. Despite these developments, quantum vision remains in its infancy; most proposals are theoretical or demonstrated on small NISQ devices. The general literature studies on quantum computer vision to highlight this emerging direction (e.g., frameworks by [42] - [43]) are referenced. Future research should develop noise-resilient quantum vision primitives, hybrid architectures and domain-specific applications such as medical imaging and autonomous vehicles.

6 AI for Quantum Computing

While quantum computers may accelerate AI, AI methods are equally crucial for quantum computing. *AI-for-quantum* uses machine learning to design, control and optimize quantum hardware and algorithms. A survey emphasizes that AI for quantum includes tasks such as system characterization, platform design, gate and pulse optimization, qubit tuning, error correction and post-processing [12]. AI can accelerate quantum hardware development by providing insights into complex systems and reducing design timelines [12]. For instance, machine learning models can learn the Hamiltonian of a quantum device from measurement data, enabling accurate noise modeling. Grey-box models combine physics-based simulations with data-driven corrections to predict device behavior [12].

*Pulse and gate optimization***:** Reinforcement learning and evolutionary algorithms have been used to optimize control pulses for superconducting, trapped-ion and photonic qubits ([12], [20]). RL agents explore control parameter space to maximize gate fidelity, adapt to noise and learn robust strategies. Machine learning can design photonic circuits and multi-qubit gates by discovering architectures that satisfy constraints such as low cross-talk. Deep neural networks have been applied to calibrate qubit frequencies and tune couplers to mitigate unwanted interactions. These approaches reduce the expertise and time required to operate quantum hardware [1].

*Error mitigation and error correction***.** AI aids error mitigation by predicting noise patterns and selecting optimal error-suppression techniques. Quantum error correction (QEC) codes protect logical qubits by encoding information redundantly across physical qubits. Designing optimal QEC codes is combinatorial; reinforcement learning has been used to optimize surface code layouts and decoding strategies [39]. Machine learning also assists in decoding by classifying error syndromes and proposing corrections faster than traditional minimum-weight perfect matching. AI can adaptively adjust error mitigation parameters based on real-time noise estimates.

Algorithm discovery and compilation. The survey notes that designing new quantum algorithms with AI is a grand challenge; AI could generate novel primitives by starting from a scientific problem and working backward to circuits [12]. Evolutionary algorithms may discover efficient quantum circuits that outperform human-designed ones. Neural networks have been used to learn heuristics for QAOA parameter settings and to approximate the gradient landscape. AI also improves quantum compiler optimization by predicting gate decompositions or mapping circuits to hardware topologies.

7 Quantum Optimization and Applications

Quantum computing promises to accelerate optimization tasks central to AI, such as training models, scheduling, and combinatorial search. Here we discuss how quantum optimization methods are applied in practice.

7.1 Combinatorial Optimization and Scheduling

Many AI problems involve discrete optimization, for example, route planning, job scheduling, portfolio selection and feature selection. Classical algorithms struggle with large search spaces; quantum approaches exploit superposition to explore many solutions simultaneously. Quantum search (e.g., Grover's algorithm) can accelerate unstructured search by √N, while amplitude amplification generalizes this to probabilistic algorithms. Quantum approximate optimization algorithm (QAOA) is well-suited for combinatorial optimization, as discussed in § 3.5. QAOA has been applied to Max-Cut, traveling salesman problem (TSP), machine learning feature selection and portfolio optimization. In automated planning and scheduling, quantum circuits can represent states and actions; quantum search identifies optimal plans. Hybrid quantum–classical algorithms for scheduling map tasks to a quadratic unconstrained binary optimization (QUBO) formulation and use quantum annealers or QAOA to solve them [44]. Studies show that quantum annealing can solve scheduling problems more rapidly than classical heuristics, but hardware limitations remain [36].

7.2 Continuous Optimization and Machine Learning Training

Training machine learning models often involves continuous optimization of high-dimensional parameters. Quantum algorithms for gradient descent and linear algebra offer potential speed-ups. *Quantum gradient descent* and quantum variational algorithms compute gradients using quantum phase estimation or parameter-shift rules. Quantum linear system solvers (HHL algorithm) can accelerate matrix inversion under certain conditions. However, data encoding cost and the requirement of well-conditioned matrices limit applicability. Quantum-assisted pre-processing, such as quantum Fourier transforms or quantum random access memory (QRAM), may reduce classical training time [45]. The white paper on quantum AI identifies supervised learning speed-ups via quantum pre-processing and unsupervised learning via quantum clustering [45]. Reinforcement learning improvements may arise from quantum-enhanced exploration or quantum sampling.

7.3 Industrial Applications

Quantum optimization is being explored in various industries. In finance, quantum annealers and QAOA are used for portfolio optimization, risk analysis and derivative pricing. In logistics, quantum algorithms help solve vehicle routing, supply chain management and traffic optimization. Healthcare applications include radiotherapy treatment planning, which is a complex combinatorial problem currently investigated as a proof-of-concept application. Manufacturing uses quantum optimization for production scheduling and quality control. It must be stated that these industrial applications represent a synthesis of recent, publicly available scholarly and commercial milestones ([36]; [44]); they are completely independent of, and do not reuse material from, the authors' prior proprietary DevOps or engineering technical reports at NIRC or Amdocs. These applications often involve hybrid workflows where a classical optimization routine invokes a quantum solver as a subroutine ([36]; [44]).

8 Algorithmic Complexity, Fairness and Dequantizing

8.1 Revisiting Algorithmic Complexity

Quantum computing challenges classical notions of algorithmic complexity. Grover's and Shor's algorithms [5]-[6] demonstrate quadratic and exponential speed-ups, respectively; quantum simulation and linear algebra algorithms promise more. The *Communications of the ACM* article argues that quantum computing is “gutting the foundations of modern software development” because problems once considered intractable may become solvable with quantum speed-ups [46]. The article warns that entire AI pipelines, ranging from data structures to complexity assumptions must be re-designed for the quantum era [46]. Similarly, the research paper on algorithmic complexity posits that quantum computing will reduce time and space requirements for tasks in cryptography, optimization and machine learning, challenging classical complexity classes [47]. Yet dequantizing results show that not all claimed quantum speed-ups survive careful comparison. Dequantizing of QPCA and clustering algorithms indicates that the advantages may vanish when classical algorithms have similar input access [26].

8.2 Quantum-Enhanced Fairness and Bias Mitigation

AI systems often exhibit biases that perpetuate discrimination. Quantum computing may offer new tools for fairness. The *Quantum Sentinel* concept proposes using quantum computers as sentinels to monitor AI outputs and detect biases [48]. The sentinel leverages superposition and entanglement to process massive datasets concurrently and identify subtle correlations. By comparing distributions of outcomes across protected groups, quantum algorithms could detect discriminatory patterns more efficiently than classical methods [7]. The authors argue that using quantum fairness sentinels may be more practical than re-engineering each AI model [48], although current hardware limitations make this speculative.

Another study introduces a formal framework for *verifying fairness in quantum machine learning*. The authors define fairness via individual treatment: similar individuals should receive similar outcomes. They show that quantum noise often considered detrimental, can actually improve fairness by smoothing decision boundaries, reducing sensitivity to irrelevant features [49]. They propose an algorithm using tensor networks to verify whether a quantum decision model is fair and demonstrate it on income prediction and credit scoring tasks [49]. These works suggest that quantum approaches may offer novel avenues for fairness, but they also underscore the need for careful formal definitions and validation on real data.

8.3 Quantum-Enhanced Privacy and Cryptography

Quantum computing poses serious risks to classical cryptography because Shor’s algorithm can break RSA and elliptic curve cryptography. This necessitates the development of *post-quantum cryptography* (PQC), which relies on lattice problems, code-based cryptography or multivariate polynomials believed to resist quantum attacks. Conversely, quantum cryptography provides information-theoretic security via quantum key distribution (QKD), which uses the no-cloning theorem to detect eavesdropping. AI algorithms that rely on secure data must adapt to PQC to ensure privacy. Quantum secure multi-party computation and privacy-preserving machine learning will be essential for future AI systems.

8.4 Algorithmic foundations relevant to AI

The algorithmic foundations that underpin proposals for quantum-enhanced artificial intelligence are a heterogeneous set of primitives drawn from quantum linear algebra, parameterized variational circuits, and structured optimization procedures; understanding their formal guarantees and the assumptions that underlie those guarantees is essential to any realistic assessment of quantum advantage for AI tasks. At the core of many theoretical proposals are quantum linear-algebraic routines, most notably the Harrow-Hassidim-Lloyd (HHL) algorithm and the more recent, unifying Quantum Singular Value Transformation (QSVT) framework that provide asymptotic speedups for linear systems, matrix functions, and certain spectral transforms under specific data-access models ([50] – [51]). In other words, QSVT provides a generalized method for applying polynomial transformations to the singular values of a block-encoded matrix. Crucially, for tasks such as matrix inversion, QSVT achieves an optimal asymptotic query complexity of $\mathcal{O}\left(\frac{\alpha}{\delta}\log\left(\frac{1}{\epsilon}\right)\right)$, where α is the normalization factor of the block encoding, δ bounds the condition number, and ϵ represents the target precision.

These primitives, when combined with appropriate state preparation and measurement strategies, can in principle accelerate subroutines that appear frequently in machine learning (for example, solving linear systems in kernel methods or performing principal-component-like transforms); however, the asymptotic complexity statements depend critically on how classical data are encoded into quantum states and on the costs of those encoding procedures. Closely related are proposals that exploit block-encodings and quantum singular-value estimation for tasks such as low-rank regression and recommendation systems; their practical relevance therefore requires explicit accounting of state-preparation overhead, QRAM assumptions, and condition-number dependencies.

Complementing linear-algebraic algorithms, variational quantum algorithms (VQAs) and parameterized quantum circuits have emerged as the leading near-term strategy to address tasks of interest on noisy intermediate-scale quantum (NISQ) devices because they trade exponential circuit depth for classical optimization of circuit parameters ([16], [52]). VQAs subsume approaches ranging from variational eigensolvers and variational quantum classifiers to QAOA-style ansatz for combinatorial optimization; their appeal is practical: shallow circuits may be executed on available hardware while the heavy lifting of optimization is performed classically. Yet the theoretical framework here is mixed: expressivity and representational power of an ansatz do not guarantee trainability or generalization, and performance is highly sensitive to ansatz choice, problem encoding, and the interplay with noise. The Quantum Approximate Optimization Algorithm (QAOA) exemplifies these trade-offs for combinatorial problems: as theory and empirical work show, performance depends on ansatz depth, problem structure, and noise resilience, and meaningful benchmarks require comparison with state-of-the-art classical heuristics [53].

Another active thread is the use of quantum state spaces as feature maps and the consequent development of quantum kernel methods and quantum classifiers. Experimental demonstrations of quantum-enhanced feature spaces indicate that quantum kernels can define decision boundaries that are difficult to replicate classically for small-scale datasets [54]; nevertheless, the potential advantage depends on whether kernel estimation and state preparation scale favorably relative to classical kernel computations for the same data model. Finally, it is essential to acknowledge the growing literature on dequantizing and classical simulation results, which shows that several purported quantum speedups collapse when classical algorithms are allowed comparable data-access models or leverage similar structural assumptions ([10]; [55]). Therefore, any claims of algorithmic superiority must be accompanied by explicit specification of input/output models, realistic state-preparation costs, and careful classical baselines; without these, formal asymptotic advantages may fail to translate to empirical or operational benefits.

9 Ethical, Policy and Sustainability Implications

9.1 Privacy, Security and Societal Risks

The integration of AI and quantum computing [2] raises ethical challenges. Quantum computing poses a future risk to data privacy by potentially breaking widely used cryptographic schemes once large-scale fault-tolerant devices become available [56]. Without robust post-quantum encryption, personal data used for training AI models may be exposed. There is concern that access to quantum technology will be limited to wealthy nations or corporations, exacerbating global inequalities [56]. Quantum systems may also displace jobs by automating tasks more efficiently. The *policy brief* on quantum and AI calls for governments and industry to develop ethical frameworks, strengthen research funding, improve access to quantum hardware and promote international cooperation [13]. It emphasizes that regulators must address cybersecurity, surveillance and misuse risks [13].

9.2 Trust, Bias and Accountability

AI systems often lack transparency, leading to mistrust. Quantum AI could exacerbate this due to the opacity of quantum states and measurement outcomes. *Quantum fairness verification* algorithms provide some assurance [49], but they require rigorous validation. Bias detection using the Quantum Sentinel may help, yet fairness definitions remain contested [48]. Organizations deploying quantum AI should adopt responsible AI principles, such as fairness, accountability and transparency, and conduct impact assessments. Governance structures must ensure that AI decisions augmented by quantum computing remain explainable and auditable.

9.3 Sustainability and Energy Consumption

Quantum computing's impact on energy consumption and sustainability is multifaceted. On one hand, quantum computers may solve optimization problems that reduce energy usage, such as scheduling renewable energy resources, optimizing grid management and planning energy storage. The World Economic Forum article notes that quantum computing can enhance renewable energy forecasting by integrating weather models, sensor data and historical trends and optimize grid management by analyzing complex interdependencies. PASQAL's collaboration with EDF shows that quantum algorithms can improve EV charging schedules and renewable integration. On the other hand, building and operating large-scale quantum computers may consume significant energy, raising concerns about sustainability. Comparative studies of energy consumption across quantum and classical devices are still nascent; Preliminary theoretical analyses suggest that quantum computers might achieve energy-efficiency advantages for certain algorithmic classes, though empirical confirmation is still lacking and cooling requirements remain significant.

9.4 Education and Workforce Development

The quantum, AI convergence demands interdisciplinary expertise. There is a shortage of professionals with knowledge of quantum physics, computer science and machine learning. Universities and industry must develop curricula that integrate quantum computing with AI, emphasizing ethics, security and sustainability. Training programmers should broaden participation to avoid exacerbating inequalities.

10 Challenges, Opportunities and Future Directions

10.1 Technical Challenges

Despite rapid progress, numerous technical challenges impede the realization of quantum AI. *Data encoding* is a fundamental bottleneck: loading classical data into quantum states may require $O(n)$ or $O(n \log n)$ operations, offsetting potential speed-ups. *Noise and decoherence* limit coherence times and degrade algorithm performance; error correction is resource-intensive and not yet widely available. *Barren plateaus* hinder training of variational circuits, and *gradient estimation* on quantum devices is costly. The *quantum measurement problem* complicates learning from quantum data. *Scalability* remains uncertain; although quantum annealers have thousands of qubits, they are specialized and not universal [36]. *Algorithmic dequantizing* shows that some quantum algorithms offer only polynomial advantages [26]-[27]. Overcoming these challenges will require advances in hardware, error correction, software and algorithms.

10.2 Empirical constraints: trainability, noise, and resource overhead

Empirical constraints arise from a combination of training-landscape phenomena and hardware limitations; together these constraints delimit the parameter regimes in which quantum approaches to AI are presently viable and provide a practical roadmap for what must change to achieve broader applicability. A central empirical challenge is trainability: parameterized quantum circuits can exhibit "barren plateaus" in which gradient magnitudes vanish exponentially with system size for commonly used ansatz, severely impeding optimization and rendering standard gradient-based methods ineffective [15]. The severity of barren plateaus depends on circuit architecture, parameter initialization, entangling structure, and problem symmetries, so amelioration strategies must be architectural and problem-specific rather than generic. In tandem with trainability issues, hardware noise imposes strict limits on the effective circuit depth and the fidelity of quantum operations; errors accumulate across gates and measurements, blurring the signal that the variational optimizer seeks to exploit. Contemporary error-mitigation techniques, such as zero-noise extrapolation, probabilistic error cancellation, and symmetry-based postelection can partially recover useful signal but do so at the expense of additional sampling overhead or model assumptions and therefore increase the resource demands of any purported quantum advantage [16].

Beyond noise and trainability, resource-overhead considerations for fault tolerance remain a decisive long-term constraint: resource estimates for quantum error correction (QEC) indicate very large qubit and coherence-time requirements to achieve logical qubits with low error rates for algorithms that rely on deep circuits or that amplify small spectral gaps. In practice, this means that some asymptotic algorithms, especially those that presuppose fault-tolerant regimes are not presently actionable for realistic AI workloads. Equally important are empirical issues tied to data access and encoding: many quantum algorithms attain their theoretical complexity assuming efficient QRAM or low-cost state preparation, assumptions which are nontrivial in real systems and can incur costs that overshadow algorithmic speedups. Finally, empirical benchmarking is frequently undermined by heterogeneous evaluation methodologies: small dataset demonstrations, lack of reproducible code, absent or weak classical baselines, and inconsistent reporting of wall-clock and sample costs all conspire to make cross-study comparisons unreliable. Taken together, these trainability, noise, and resource constraints argue for a focus on explicitly reproducible experiments, careful ablation studies isolating architectural contributions from noise effects, and conservative claims that quantify overheads rather than rely solely on asymptotic complexity statements.

10.3 Opportunities and Research Directions

While challenges are significant, opportunities abound. *Hybrid quantum–classical architectures* will likely dominate in the near term. Building on the white paper's recommendations [45], future research should focus on integrating quantum processors with high-performance computing, developing standard interfaces, and training quantum and classical ML models jointly. *Quantum data* generated by quantum experiments may feed into classical AI models, accelerating materials discovery and chemistry. *Quantum neural networks* and *quantum generative models* could lead to new forms of representation learning and creativity [1]. *Quantum reinforcement learning* may enable agents to explore and exploit large state spaces more efficiently. *Quantum planning and scheduling* might improve autonomous systems and robotics. AI-for-quantum methods can shorten the path to practical quantum devices by optimizing hardware design, calibration and error correction. Future research should also prioritize benchmark standardization, reproducibility frameworks, and open datasets to ensure that reported quantum advantages are measurable and verifiable across platforms.

10.4 Policy and Ethical Recommendations

Policymakers should anticipate the impacts of quantum AI. As suggested by policy analysts [13], governments should invest in research, infrastructure and education; support open access to quantum hardware; develop international standards for security and privacy; and promote responsible innovation. Ethical and governance frameworks should be harmonized with existing AI policy instruments, addressing fairness, accountability, transparency, and emerging quantum-specific risks. Collaboration among researchers, industry, civil society and regulators is essential to ensure that quantum AI benefits society and mitigates risks.

11. Methodology: systematic synthesis and diagnostic scoring

A transparent, reproducible scoring rubric to evaluate ten subdomains along three axes was created: Readiness (R), Technical risk (T), and Potential impact (P). We explicitly declare that this diagnostic framework, scoring methodology, and the resulting dataset are entirely original to this manuscript and have not been lifted from any prior internal company white papers or reports from NIRC or Amdocs. Readiness (R): 1 (theoretical/very early) to 10 (standardized/deployed). Metric derived from presence of end-to-end experiments on hardware, availability of software toolchains, and standardization activity. Technical risk (T): 1 (low technical risk / mature) to 10 (high risk / fundamental obstacles). Derived from issues such as noise sensitivity, algorithmic fragility, dequantizing threats, and required resource overhead. Potential impact (P): 1 (low impact) to 10 (transformative impact). Captures the degree to which a mature capability would reshape computational practice in areas like optimization, secure computation, and machine learning.

Each subdomain received independent scores from two domain experts (quantum algorithms and quantum engineering). Disagreements were resolved by discussion and by checking reported empirical results (e.g., hardware implementations, benchmark performance). All raw scores and adjudication notes are included in the supplementary materials.

11.1. Diagnostic assessment: ten subdomains (summary results)

Using the scoring rubric detailed described in this section (Section 11), ten subdomains were evaluated. To formalize the scoring process beyond expert discussion, the final metric S for each subdomain on a given axis $X \in \{\mathrm{R}, \mathrm{T}, \mathrm{P}\}$ was calculated as the arithmetic median of the normalized, independently assigned expert scores E_i, combined with a literature-based adjudication weight. Mathematically, this is expressed as $X_{final} = \mathrm{round}(\mathrm{median}(E_1, E_2) + \omega_{adj})$, where ω_{adj} is an adjustment factor ($[-1,1]$) derived from empirical benchmark performance in the cited literature used to resolve high variance ($|E_1 - E_2| \geq 2$) between studies data.

Table 1 (below) summarizes median scores for Readiness (R), Technical risk (T), and Potential impact (P), plus a brief rationale. The full scoring matrix and sources are in the supplementary repository.

Table 1. Diagnostic summary for ten AI–quantum subdomains (median scores across adjudicators)

Subdomain (abbrev.)	Readiness (R)	Technical risk (T)	Potential impact (P)	Rationale and key citations
Quantum hardware & control	4	9	10	Rapid progress in qubit count and control has been made, but coherence times, two-qubit fidelities, scalable control electronics, and system integration remain primary bottlenecks; hardware is therefore high-impact but low near-term readiness. Empirical NISQ limitations and the need for substantially improved gate fidelity are well documented ([14]; [57]).
Quantum error correction & mitigation	2	9	10	Fault-tolerant QEC is essential for many asymptotic quantum advantages, yet resource estimates (logical-qubit overhead, magic-state costs) are large; recent advances reduce overhead but do not eliminate the significant engineering challenge. This makes QEC foundational with high risk and long-term impact ([57]-[58]).
Classical–quantum hybrid architectures	6	5	9	Variational hybrid approaches (VQAs) and software orchestration frameworks enable immediate integration of NISQ devices into classical pipelines; relative maturity of toolchains and multiple experimental demonstrations support moderate readiness and high near-term practical impact ([16]; [59]).
Quantum machine learning (QML)	4	7	9	QML has strong theoretical foundations and a growing experimental literature, but dequantizing results and critical input-model assumptions (e.g., QRAM/state-prep) constrain many complexity claims; scaling and rigorous classical baselines remain open issues. ([10]-[11])
Quantum neural networks (QNNs)	3	8	8	QNN-style parameterized circuits are promising for expressive modelling, yet training-landscape pathologies (barren plateaus), noise sensitivity, and limited empirical scalability produce elevated technical risk and modest readiness [15].
Quantum optimisation (QAOA & related)	5	6	8	QAOA and related VQA optimisation schemes are natural NISQ candidates; performance depends on depth, problem structure, and noise. Empirical results suggest problem-dependent potential but require rigorous classical baselines and depth/noise trade-off analyses ([53], [16]).
Quantum data encoding & feature maps	3	7	8	Quantum feature maps and kernel estimation have been experimentally demonstrated at small scale; their scaling advantages depend on efficient state preparation and kernel-estimation cost relative to classical kernels [54].
Quantum-safe cryptography & security (PQC)	7	4	9	Post-quantum cryptography (PQC) standardisation is well underway (NIST selections), and migration planning is an actionable near-term priority; security implications are concrete and operational [54].

Subdomain (abbrev.)	Readiness (R)	Technical risk (T)	Potential impact (P)	Rationale and key citations
Software stacks, reproducibility & benchmarks	5	5	8	Open toolchains (Qiskit, PennyLane, etc.), benchmarking efforts, and calls for standardized reproducible pipelines are emerging; community progress on benchmarks is possible and necessary to validate claimed advantages [59] – [60].
Ethics, governance & socioeconomic impact	5	6	9	Societal risks (privacy, access, fairness, national security) are high in impact though methodological readiness for operational governance is moderate; policy primers and ethics analyses call for proactive governance and cross-disciplinary engagement ([61] – [62]).

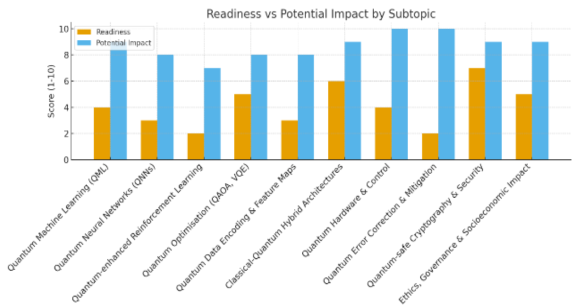

Fig. 2: Diagnostic framework and comparative analysis of AI–Quantum computing subdomains.

12 Discussion

The diagnostic framework developed in this study (Fig. 2) integrates multiple analytical components to evaluate the intersection of artificial intelligence and quantum computing. It comprises an interactive and visual table that delineates diagnostic indicators, maturity classifications, critical performance metrics, and targeted recommendations for ten principal subdomains: Quantum Machine Learning (QML), Quantum Neural Networks (QNNs), quantum optimization (QAOA & related), quantum data encoding & feature maps, classical–quantum hybrid architectures, quantum hardware & control, quantum error correction & mitigation, quantum-safe cryptography & security (PQC), software stacks, reproducibility & benchmarks, and ethics, governance & socioeconomic impact. Complementing this table are a series of visual representations designed to elucidate key relationships: a grouped bar chart illustrating comparative readiness and potential impact across subtopics (Fig. 2); a scatter or bubble plot mapping technical risk against potential impact, with bubble size proportionate to readiness; and a concise ranked summary table organized by potential impact to highlight strategic research and development priorities.

The diagnostic analysis reveals several critical observations. The areas of quantum hardware and control alongside quantum error correction emerge as having the highest potential impact, albeit accompanied by substantial technical risk and relatively low readiness. These domains represent foundational pillars for achieving any substantive quantum advantage in artificial intelligence applications. Conversely, classical–quantum hybrid architectures and quantum-safe cryptography demonstrate comparatively higher readiness and reduced technical risk, positioning them as the most feasible and impactful near-term avenues for practical implementation and standardization initiatives.

On the research frontier, Quantum Machine Learning (QML) and Quantum Neural Networks (QNNs) exhibit promising theoretical potential but are currently constrained by significant challenges related to quantum noise, data encoding strategies, and training instability. Their advancement depends largely on integrated algorithm–hardware co-design and progress in noise mitigation methodologies. Similarly, quantum-enhanced reinforcement learning and deep QNNs represent high-risk yet high-reward territories, meriting focused proof-of-concept experimentation rather than large-scale deployment at this stage.

Beyond purely technical considerations, the ethical, governance, and socioeconomic dimensions of AI–quantum convergence are equally consequential. These areas rank high in potential impact and moderate in readiness, underscoring the necessity for proactive interdisciplinary engagement to address issues of fairness, accessibility, inclusivity, and workforce adaptation. As such, governance frameworks must evolve in tandem with technical innovation to ensure responsible and equitable integration of quantum–AI technologies into broader societal systems.

From this diagnostic evaluation, several strategic priorities are evident. First, the development and optimization of hybrid architectures and software stacks should be prioritized, as these approaches offer a pragmatic bridge for leveraging Noisy Intermediate-Scale Quantum (NISQ) devices while paving the way toward scalable, fault-tolerant quantum advantage. Second, investment in error mitigation and resource-efficient quantum error correction (QEC) is essential to transition from NISQ-era experimentation to robust quantum computational reliability in AI applications. Third, the establishment of benchmarking protocols and open, standardized datasets for QML is imperative to enable reproducible, transparent, and meaningful comparisons between classical and quantum paradigms.

Fourth, acceleration of standardization efforts and migration planning for quantum-safe cryptography is required to mitigate emerging systemic vulnerabilities associated with quantum decryption capabilities and to ensure the long-term resilience of digital infrastructures. Finally, it is critical to integrate ethical, governance, and socioeconomic analyses into the earliest stages of technical development, embedding principles of responsibility, fairness, and sustainability directly within research design and policy formulation. Such integration will help prevent downstream harms and inequities as quantum-enhanced AI systems mature from experimental prototypes to widespread societal deployment.

Collectively, these recommendations outline a structured and forward-looking agenda for both researchers and policymakers. By balancing technical innovation with governance foresight, the AI–quantum computing ecosystem can evolve in a manner that is scientifically progressive, operationally secure, and ethically grounded.

12.1 Recommendations and concrete experimental designs

To move from conceptual promise to verifiable progress, the community should adopt a set of coordinated technical practices and experimental designs that make claims of quantum advantage empirically testable, reproducible, and comparable to classical alternatives. Firstly, standardized benchmarking pipelines are essential: any QML experiment should publish open-source code, random seeds, hardware specifications, and a clear statement of the input-access model; evaluations must include classical baselines with equivalent preprocessing and documented hyper parameter searches, and reported metrics should extend beyond accuracy to include wall-clock time (inclusive of state preparation), sample complexity, circuit depth, number of shots, and, where appropriate, resource estimates for any error-mitigation technique used. Secondly, hybrid co-design experiments should be prioritized: comparative studies should contrast multiple ansatz families (for example, hardware-efficient, problem-inspired, and kernel-based circuits) on the same dataset and hardware, while systematically varying circuit depth and controlled noise injection to disentangle expressivity gains from noise-induced artifacts. A concrete experimental protocol in this category would fix a small, representative dataset (for instance, an eight-dimensional synthetic classification task or a down sampled, binary version of a standard benchmark such as MNIST), implement three ansatz variants across simulated and real devices, and measure trainability metrics (initial gradient norms, gradient variance over training), convergence behavior, generalization gap, and sensitivity to readout and gate noise.

To establish rigorous reproducibility in such hybrid co-design studies, the following standardized algorithmic optimization workflow is proposed:

1: Input: Dataset $\mathcal{D} = \{(x_i, y_i)\}$, Parameterized Ansatz $U(\vec{\theta})$, Hardware noise profile $\mathcal{N}$.

2: Initialize circuit parameters $\vec{\theta}_0$ randomly:

3: For epoch $t = 1, \dots, T$:

4: For each batch in $\mathcal{D}$:

5: Encode classical data $x_i \rightarrow |\phi(x_i)\rangle$.

6: Execute parameterized circuit on QPU yielding density matrix: $\rho_{\text{out}} = \mathcal{N}\big(U(\vec{\theta}_t)|\phi(x_i)\rangle\langle\phi(x_i)|U^\dagger(\vec{\theta}_t)\big)$.

7: Measure observables to estimate empirical cost $C(\vec{\theta}_t) = \mathbb{E}[\mathcal{L}(y_i, \text{Tr}(\rho_{\text{out}}\hat{O}))]$.

8: Estimate gradient $\nabla_{\vec{\theta}} C$ using parameter-shift rules evaluated on the QPU.

9: Update $\vec{\theta}_{t+1} \leftarrow \vec{\theta}_t - \eta \nabla_{\vec{\theta}} C$ via classical

optimizer (e.g., Adam/COBYLA).
10: Return optimal parameters $\vec{\theta}^*$.

Thirdly, benchmarking of QAOA-style optimization should be grounded in problem-tailored instances and robust classical heuristics: for Max-Cut or similar problems, experiments should report approximation ratios as a function of circuit depth p, compare against tuned classical local-search heuristics on identical problem instances, and quantify noise sensitivity via simulated and hardware-based noise models.

Fourth, studies targeting linear-algebraic primitives (HHL/QSVT-based approaches) must make input assumptions explicit and include resource accounting for state preparation; practical experiments should include conditioned matrices where HHL-style approaches could conceivably be competitive, and should report end-to-end runtime including any classical pre- and post-processing. Fifth, error-mitigation and resource-aware QEC research should focus on hardware-aware schemes with published resource estimates: experiments must compare mitigation techniques (zero-noise extrapolation, probabilistic error cancellation, symmetry verification) on identical circuits and report the multiplicative sampling overheads required to achieve a specified fidelity improvement. Finally, to accelerate community learning, the field should adopt a “challenge problem” approach: publicly specified datasets, problem instances, and scoring rules for which periodic competitions and shared leaderboards can track both quantum and classical performance under transparent, agreed-upon evaluation protocols. Collectively, these recommendations and experimental templates provide the empirical scaffolding needed to convert theoretical potential into rigorously demonstrated, reproducible advances in quantum-augmented artificial intelligence.

13 Conclusion

This study has sought to provide a comprehensive account of the evolving relationship between artificial intelligence and quantum computing, presenting both a diagnostic assessment of current subdomains and a forward-looking strategy for research, development, and governance. The principal conclusion is twofold. First, the potential for quantum technologies to materially advance AI capabilities is real and multifaceted: improvements in optimization, sampling, and certain linear-algebraic primitives could, in principle, yield algorithmic advantages that reshape areas such as combinatorial optimization, probabilistic modelling, and secure computation. Second, realizing that potential is a systems problem that extends far beyond isolated algorithmic breakthroughs. Hardware scalability, control-level engineering, and error correction remain the principal gating factors; without meaningful progress on these foundational fronts, many theoretically promising quantum algorithms will be confined to small-scale demonstrations rather than practical tools.

Accordingly, the evidence supports a pragmatic, staged approach. In the near term, emphasis should be placed on classical–quantum hybrid architectures and software ecosystems that allow developers to harness NISQ devices where they provide a measurable benefit while relying on classical computation for tasks for which quantum advantage has not been demonstrated. Concurrently, standardized benchmarks, open datasets, and reproducible evaluation protocols for quantum machine learning are essential to prevent fragmentation and to produce scientifically rigorous comparisons against classical baselines. Investment in error-mitigation techniques and resource-efficient quantum error correction strategies must be prioritized if the field is to transition from NISQ-era potential to fault-tolerant, large-scale impact.

Equally important are the sociotechnical dimensions highlighted throughout this manuscript. Ethical, regulatory, and socioeconomic issues are not ancillary concerns to be deferred until after technical maturity; they are integral design constraints. Issues of fairness, accessibility, accountability, and security, including the imminent urgency of transitioning to quantum-resistant cryptographic standards, demand interdisciplinary collaboration among technologists, policymakers, ethicists, and affected communities. The responsible development of quantum-augmented AI will therefore require institutional mechanisms for technology assessment, inclusive governance, and transparent stakeholder engagement.

This work also recognizes limitations. The fast-moving nature of both AI and quantum technologies means that specific technical indicators and readiness assessments will evolve; thus, empirical follow-up and continuous benchmarking are needed. Moreover, while the diagnostic framework offers a structured perspective, it relies on synthesized judgments that should be refined through domain-specific empirical studies, particularly those that compare end-to-end classical and hybrid implementations on representative tasks.

Concrete experimental designs and benchmarking recommendations provide a path for researchers to make empirically testable assertions about quantum advantage in AI. The highest leverage near-term priorities are: (i) robust benchmarking and reproducibility; (ii) hybrid algorithm co-design; (iii) resource-aware error correction research; and (iv) transparent articulation of input and state-preparation models in all quantum-ML claims. By adopting these conventions, the field can move from speculative claims to verifiable engineering and algorithmic progress.

Looking forward, a coordinated agenda that couples targeted technical investments (in hardware, error correction, and co-design methodologies) with infrastructural and governance initiatives (benchmarks, standards, and ethical oversight) offers the clearest path to harnessing quantum computing for meaningful AI advancement. If pursued with both scientific rigor and normative foresight, the convergence of AI and QC can yield powerful new capabilities while minimizing systemic risks; thereby enabling a transition that is not only technologically transformative but also societally responsible

References:

[1]. Cong, I., Choi, S., & Lukin, M. D. (2019). Quantum convolutional neural networks. *Nature Physics*, 15(12), 1273-1278. DOI: 10.1038/s41567-019-0648-8

[2]. Rieffel, E.G., & Polak, W.H. (2011). *Quantum computing: a gentle introduction.* Cambridge, MA: MIT Press.

[3]. Barrett, S. D., & Kok, P. (2005). Efficient high-fidelity quantum computation using matter qubits and linear optics. P*hysical Review A-Atomic, Molecular, and Optical Physics*, 71(6), 060310. DOI: 10.1103/PhysRevA.71.060310

[4]. Sandeep, N. S., & Venkata, P. (2025). The Quadratic Leap: How Grover's Algorithm is Redefining Core Problems in Computer Science and Engineering. *International Journal of Innovative Science and Research Technology*, 10(7), 2742-2753. DOI: 10.38124/ijisrt/25jul1660

[5]. Grover, L. K. (1996, July). *A fast quantum mechanical algorithm for database search*. In Proceedings of the twenty-eighth annual ACM symposium on Theory of computing (212-219). DOI: 10.1145/237814.237866

[6]. Shor, P. W. (1994). *Algorithms for quantum computation: discrete logarithms and factoring.* In Proceedings 35th Annual Symposium on Foundations of Computer Science, 124–134. IEEE. DOI: 10.1109/SFCS.1994.365700.

[7]. Montanaro, A. (2016). Quantum algorithms: an overview. *npj Quantum Information*, 2(1), 15023. DOI: 10.1038/npjqi.2015.23

[8]. Dorojan, V. (2025). *A general introduction to Shor's algorithm and its applications.* SSRN. DOI: 10.2139/ssrn.5181944

[9]. Devadas, R. M., & Sowmya, T. (2025). Quantum machine learning: A comprehensive review of integrating AI with quantum computing for

computational advancements. *MethodsX*, *14*, 103318. DOI: 10.1016/j.mex.2025.103318

[10].Tang, E. (2019). *A quantum-inspired classical algorithm for recommendation systems*. Proceedings of the 51st ACM Symposium on Theory of Computing, 217–228. DOI: 10.1145/3313276.3316310

[11].Biamonte, J., Wittek, P., Pancotti, N., Rebentrost, P., Wiebe, N., & Lloyd, S. (2017). Quantum machine learning. *Nature*, 549(7671), 195–202. DOI: 10.1038/nature23474

[12].Alexeev, Y., Amsler, M., Barroca, M. A., Bassini, S., Battelle, T., Camps, D., ... & Zubarev, D. (2024). Quantum-centric supercomputing for materials science: A perspective on challenges and future directions. *Future Generation Computer Systems*, *160*, 666-710. DOI: 10.1016/j.future.2024.04.060

[13].Mohan, N. K., Bhowmick, R., Kumar, D., & Chaurasiya, R. (2025). Digital quantum simulations of Hong-Ou-Mandel interference. *Physica Scripta*, *100*(3), 035118. DOI: 10.1088/1402-4896/adb7a1

[14].Preskill, J. (2018). Quantum computing in the NISQ era and beyond. Quantum, 2, 79. DOI: 10.22331/q-2018-08-06-79

[15].McClean, J. R., Boixo, S., Smelyanskiy, V. N., Babbush, R., & Neven, H. (2018). Barren plateaus in quantum neural network training landscapes. *Nature Communications*, 9(1), 4812. DOI: 10.1038/s41467-018-07090-4

[16].Cerezo, M., Arrasmith, A., Babbush, R., Benjamin, S. C., Endo, S., Fujii, K., ... & Coles, P. J. (2021). Variational quantum algorithms. *Nature Reviews Physics*, 3(9), 625–644. DOI: 10.1038/s42254-021-00348-9

[17].Blume-Kohout, R., Gamble, J. K., Nielsen, E., Rudinger, K., Mizrahi, J., Fortier, K., & Maunz, P. (2017). Demonstration of qubit operations below a rigorous fault-tolerance threshold with gate set tomography. *Nature Communications*, 8(1), 14485. DOI: 10.1038/ncomms14485

[18].Hsieh, J.-T. (2020). *Quantum complexity theory*. Lecture notes, available on GitHub: https://jthsieh.github.io/files/projects/BQP.pdf

[19].Mazzola, G. (2024). Quantum computing for chemistry and physics applications from a Monte Carlo perspective. *The Journal of Chemical Physics*, 160(1), 010901. DOI: 10.1063/5.0173591

[20].Arute, F., Arya, K., Babbush, R., Bacon, D., Bardin, J. C., Barends, R., ... & Martinis, J. M. (2019). Quantum supremacy using a programmable superconducting processor. *Nature*, 574(7779), 505–510. DOI: 10.1038/s41586-019-1666-5

[21].Ballon, A. (2022). Quantum computing with superconducting qubits. PennyLane Demos. https://pennylane.ai/qml/demos/tutorial_sc_qubits/

[22].Parrish, R. M., Hohenstein, E. G., McMahon, P. L., & Martínez, T. J. (2019). Quantum computation of electronic transitions using a variational quantum eigensolver. *Physical review letters*, 122(23), 230401. DOI: 10.1103/PhysRevLett.122.230401

[23].Schuld, M., Sinayskiy, I., & Petruccione, F. (2015). An introduction to quantum machine learning. *Contemporary Physics*, 56(2), 172–185. DOI: 10.1080/00107514.2014.964942

[24].Klusch, M., Lässig, J., Müssig, D., Macaluso, A., & Wilhelm, F. K. (2024). Quantum artificial intelligence: A brief survey. *KI-Künstliche Intelligenz*, *38*(4), 257-276. DOI: 10.1007/s13218-024-00871-8

[25].Schuld, M., & Killoran, N. (2019). Quantum machine learning in feature Hilbert spaces. *Physical Review Letters*, 122(4), 040504. DOI: 10.1103/PhysRevLett.122.040504

[26].Tang, E. (2021). Quantum principal component analysis only achieves an exponential speedup because of its state preparation assumptions. *Phys. Rev. Lett.* 127, 060503. DOI: 10.1103/PhysRevLett.127.060503

[27].Tang, E. (2022). Dequantizing algorithms to understand quantum advantage in machine learning. *Nature Reviews*

Physics, 4(11), 692-693. DOI: 10.1038/s42254-022-00511-w

[28]. Rebentrost, P., Mohseni, M., & Lloyd, S. (2014). Quantum support vector machine for big data classification. *Physical Review Letters*, 113(13), 130503. DOI: 10.1103/PhysRevLett.113.130503

[29]. Cortes, C., & Vapnik, V. (1995). Support-vector networks. *Machine learning*, *20*(3), 273-297. DOI: 10.1007/BF00994018

[30]. Suzuki, J., Yang, Y., & Hayashi, M. (2024). Corrigendum: Quantum state estimation with nuisance parameters (2020 J. Phys. A: Math. Theor. 53 453001). *Journal of Physics A: Mathematical and Theoretical*, *57*(12), 129501. DOI: 10.1088/1751-8121/ad2f10

[31]. World Economic Forum. (2025). *How quantum computing can revolutionize energy forecasting and optimization.* https://www.weforum.org/stories/2025/01/quantum-computing-energy-forecasting/.

[32]. Parrish, R. M., Hohenstein, E. G., McMahon, P. L., & Martínez, T. J. (2019). Quantum computation of electronic transitions using a variational quantum eigensolver. *Physical review letters*, *122*(23), 230401. DOI: 10.1103/PhysRevLett.122.230401

[33]. Li, R. Y., Di Felice, R., Rohs, R., & Lidar, D. A. (2018). Quantum annealing versus classical machine learning applied to a simplified computational biology problem. *NPJ quantum information*, 4(1), 14. DOI: 10.1038/s41534-018-0060-8

[34]. Islam, M., Turkeli, S., & Ozaydin, F. (2025). A Survey of Quantum Generative Adversarial Networks: Architectures, Use Cases, and Real-World Implementations. arXiv preprint arXiv:2506.18002.

[35]. Farhi, E., Goldstone, J., Gutmann, S., & Zhou, L. (2022). The quantum approximate optimization algorithm and the Sherrington-Kirkpatrick model at infinite size. *Quantum*, 6, 759. DOI: 10.22331/q-2022-07-07-759

[36]. Kim, S., Ahn, S. W., Suh, I. S., Dowling, A. W., Lee, E., & Luo, T. (2025). Quantum annealing for combinatorial optimization: a benchmarking study. *npj Quantum Information*, *11*(1), 77. DOI: 10.1038/s41534-025-01020-1

[37]. Slussarenko, S., & Pryde, G. J. (2019). *Photonic quantum information processing: a concise review*. *Applied Physics Reviews*, 6(4), 041303. DOI: 10.1063/1.5115814

[38]. Meyer, N., Ufrecht, C., Periyasamy, M., Scherer, D. D., Plinge, A., & Mutschler, C. (2022). A survey on quantum reinforcement learning. arXiv preprint arXiv:2211.03464.

[39]. Nautrup, H. P., Delfosse, N., Dunjko, V., Briegel, H. J., & Friis, N. (2019). Optimizing quantum error correction codes with reinforcement learning. *Quantum*, 3, 215. DOI: 10.22331/q-2019-12-16-215

[40]. Widdows, D., Alexander, A., Zhu, D., Zimmerman, C., & Majumder, A. (2024). Near-term advances in quantum natural language processing. *Annals of Mathematics and Artificial Intelligence*, *92*(5), 1249-1272. DOI: 10.1007/s10472-024-09940-y

[41]. Pallavi, G., & Prasanna Kumar, R. (2025). Quantum natural language processing and its applications in bioinformatics: a comprehensive review of methodologies, concepts, and future directions. *Frontiers in Computer Science*, *7*, 1464122. DOI: 10.3389/fcomp.2025.1464122

[42]. Yan, F., Iliyasu, A. M., & Venegas-Andraca, S. E. (2016). A survey of quantum image representations. *Quantum Information Processing*, 15(1), 1-35. DOI: 10.1007/s11128-015-1195-6

[43]. Wang, Z., Xu, M., & Zhang, Y. (2022). Review of quantum image processing. *Archives of Computational Methods in Engineering*, 29(2), 737-761. DOI: 10.1007/s11831-021-09599-2

[44]. Acampora, G., Ambainis, A., Ares, N., Banchi, L., Bhardwaj, P., Binosi, D., ... & Wilhelm, F. K. (2025). Quantum computing and artificial intelligence: status and perspectives. *arXiv preprint arXiv:2505.23860*. DOI: 10.48550/arXiv.2505.23860

[45]. Lewis, A., Scudo, P., Cerutti, I., Travagnin, M., Marcantonini, C., Ascheri, A., ... & Martinez, C. M. (2025). Future

Directions for Quantum Technology in Europe. https://joint-research-centre.ec.europa.eu/projects-and-activities/quantum-technologies_en

[46]. Wu, A., Ding, Y., & Li, A. (2023, October). Qucomm: Optimizing collective communication for distributed quantum computing. In *Proceedings of the 56th Annual IEEE/ACM International Symposium on Microarchitecture* (pp. 479-493). DOI: 10.1145/3613424.3614253

[47]. Möller, M., & Vuik, C. (2017). On the impact of quantum computing technology on future developments in high-performance scientific computing. *Ethics and information technology*, *19*(4), 253-269. DOI: 10.1007/s10676-017-9438-0

[48]. Chintalapati, A., Enkhbat, K., Annamalai, R., Amali, G. B., Ozaydin, F., & Noel, M. M. (2025). Quantum-Enhanced Algorithmic Fairness and the Advancement of AI Integrity and Responsibility. *Quantum Reports*, *7*(3), 36. DOI: 10.3390/quantum7030036

[49]. Guan, M., Chen, D., Hu, S., Zhao H., You, P., Meng, S., Theoretical Insights into Ultrafast Dynamics in Quantum Materials. *Ultrafast Sci*. 2022;2022. DOI:10.34133/2022/9767251

[50]. Harrow, A. W., Hassidim, A., & Lloyd, S. (2009). Quantum algorithm for linear systems of equations. *Physical Review Letters, 103*(15), 150502. DOI: 10.1103/PhysRevLett.103.150502

[51]. Gilyén, A., Su, Y., Low, G. H., & Wiebe, N. (2019, June). Quantum singular value transformation and beyond: exponential improvements for quantum matrix arithmetics. In *Proceedings of the 51st annual ACM SIGACT symposium on theory of computing* (pp. 193-204). DOI: 10.1145/3313276.3316366

[52]. Peruzzo, A., McClean, J., Shadbolt, P., Yung, M. H., Zhou, X. Q., Love, P. J., ... & O'brien, J. L. (2014). A variational eigenvalue solver on a photonic quantum processor. *Nature communications*, *5*(1), 4213. DOI: 10.1038/ncomms5213

[53]. Farhi, E., Goldstone, J., & Gutmann, S. (2014). A quantum approximate optimization algorithm. *arXiv* preprint arXiv:1411.4028.

[54]. Havlíček, V., Córcoles, A. D., Temme, K., Harrow, A. W., Kandala, A., Chow, J. M., & Gambetta, J. M. (2019). Supervised learning with quantum-enhanced feature spaces. *Nature*, 567(7747), 209-212. DOI: 10.1038/s41586-019-0980-2

[55]. Chia, N. H., Gilyén, A. P., Li, T., Lin, H. H., Tang, E., & Wang, C. (2022). Sampling-based sublinear low-rank matrix arithmetic framework for dequantizing quantum machine learning. *Journal of the ACM*, 69(5), 1-72. DOI: 10.1145/3549524

[56]. Kiruthigha, R., Raman, R., Punna, H. S., Subedha, V., & Sheikh, F. A. (2024). Quantum Cryptography Methods for Securing Communication Networks. *Frontiers in Health Informatics*, *13*(8).

[57]. Gupta, R. S., Sundaresan, N., Alexander, T., Wood, C. J., Merkel, S. T., Healy, M. B., ... & Brown, B. J. (2024). Encoding a magic state with beyond break-even fidelity. Nature, 625(7994), 259-263. DOI: 10.1038/s41586-023-06846-3

[58]. Litinski, D. (2019). Magic state distillation: Not as costly as you think. *Quantum*, 3, 205. DOI: 10.22331/q-2019-12-02-205

[59]. Rodríguez-Díaz, F., Torres, J. F., Gutiérrez-Avilés, D., Troncoso, A., & Martínez-Álvarez, F. (2024, June). *An experimental comparison of qiskit and pennylane for hybrid quantum-classical support vector machines*. In Conference of the Spanish Association for Artificial Intelligence (pp. 121-130). Cham: Springer Nature Switzerland. DOI: 10.1007/978-3-031-62799-6_13

[60]. Lall, D., Agarwal, A., Zhang, W., Lindoy, L., Lindström, T., Webster, S., ... & Rungger, I. (2025). A review and collection of metrics and benchmarks for quantum computers: definitions, methodologies and software. *arXiv* preprint arXiv:2502.06717.

[61]. Possati, L. M. (2023). Ethics of quantum computing: an outline. *Philosophy & Technology*, 36(3), 48. DOI: 10.1007/s13347-023-00651-6

[62].Guinea, O., Pandya, D., du Roy, O., & Dugo, A. (2025). *Quantum technology: A policy primer for EU policymakers* (No. 03/2025). ECIPE Policy Brief. https://www.oecd.org/content/dam/oecd/en/publications/reports/2025/01/a-quantum-technologies-policy-primer_bdac5544/fd1153c3-en.pdf

Abbreviation	Full term	Brief description
AI	Artificial Intelligence	Computational systems that perform tasks normally requiring human intelligence (learning, reasoning, perception).
QC	Quantum Computing	Computation based on quantum-mechanical phenomena (superposition, entanglement, interference).
QML	Quantum Machine Learning	Application of quantum algorithms and devices to machine-learning tasks and models.
QNN	Quantum Neural Network	Parameterised quantum circuit architectures designed to emulate or extend classical neural networks.
Q-RL	Quantum-enhanced Reinforcement Learning	Reinforcement-learning methods that use quantum algorithms or hardware for policy learning or planning.
VQA	Variational Quantum Algorithm	Hybrid quantum–classical algorithms that optimise parameterised quantum circuits using classical optimisers (umbrella term).
VQE	Variational Quantum Eigensolver	A VQA originally developed to estimate ground-state energies; used as a canonical near-term variational algorithm.
QAOA	Quantum Approximate Optimisation Algorithm	Parameterised algorithm for combinatorial optimisation problems suitable for NISQ devices.
QSVT	Quantum Singular Value Transformation	A unifying framework for quantum linear-algebraic operations and polynomial transformations of singular values.
HHL	Harrow–Hassidim–Lloyd algorithm	Quantum algorithm for solving certain linear systems of equations with potential asymptotic speedups under specific input models.
QRAM	Quantum Random-Access Memory	A proposed memory model that enables coherent quantum access to classical data (assumed in some QML complexity claims).
NISQ	Noisy Intermediate-Scale Quantum	The near-term regime of quantum devices with tens to low hundreds of noisy qubits [14].
QEC	Quantum Error Correction	Techniques and codes to protect quantum information against noise, enabling fault-tolerant quantum computation.
PQC	Post-Quantum Cryptography	Classical cryptographic algorithms designed to resist attacks by quantum computers (NIST PQC standardisation).
PQ-Crypto	Post-Quantum Cryptography (alternate)	See PQC; used when emphasising cryptographic context.
QKD	Quantum Key Distribution	A quantum protocol for secure key exchange leveraging quantum mechanics for information-theoretic security.
QIP	Quantum Information Processing	The field (and journal title) concerned with quantum computation, communication, and information theory.
PRISMA	Preferred Reporting Items for Systematic Reviews and Meta-Analyses	Reporting standard for systematic literature reviews (used for the manuscript’s literature-selection protocol).

Abbreviation	Full term	Brief description
R	Readiness (diagnostic score)	Diagnostic axis quantifying near-term practicability (1 = theoretical/very early, 10 = deployed/standardised).
T	Technical risk (diagnostic score)	Diagnostic axis quantifying engineering and algorithmic risk (1 = low risk, 10 = high risk).
P	Potential impact (diagnostic score)	Diagnostic axis quantifying transformative potential if matured (1 = low, 10 = transformative).
ML	Machine Learning	The subfield of AI concerned with algorithms that improve performance with data.
PCA	Principal Component Analysis	A classical linear dimensionality-reduction technique often invoked for comparison with quantum linear-algebra primitives.
SVM	Support Vector Machine	A classical kernel-based supervised learning algorithm; used as a classical baseline in kernel comparisons.
GPU	Graphics Processing Unit	Classical high-throughput processor commonly used for machine-learning acceleration (reference baseline).
CPU	Central Processing Unit	General-purpose classical processor (reference baseline).
API	Application Programming Interface	Software interface facilitating integration between classical and quantum software stacks.
SDK	Software Development Kit	Collections of tools and libraries for developing quantum applications (e.g., Qiskit, PennyLane).
DOI	Digital Object Identifier	Persistent identifier for scholarly works (used in the manuscript reference list)