Exponential Science
Discussion Paper Series
No 13-25

A Classical Perspective on the Synergies of Quantum and Artificial Intelligence

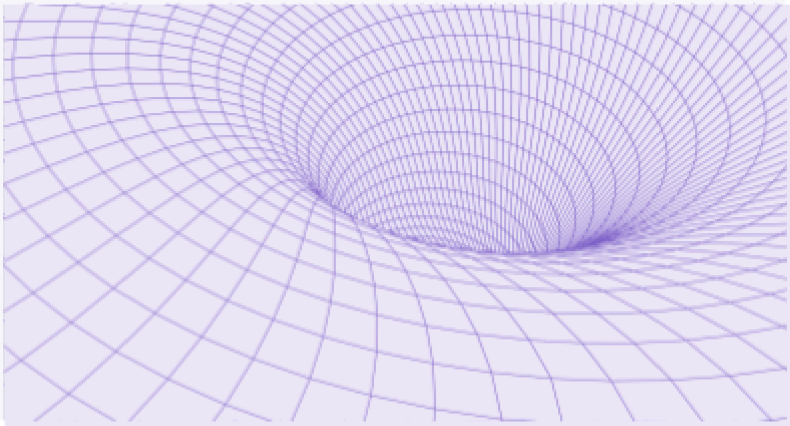

Hala Hawashin
Exponential Science
University College London

Marco Alberto Javarone
Exponential Science

Hala Hawashin1, Marco Alberto Javarone. A Classical Perspective on the Synergies of Quantum and Artificial Intelligence. Printed in ES Discussion Paper Series (Exponential Science, 2025).

A Classical Perspective on the Synergies of Quantum and Artificial Intelligence

Hala Hawashin[1] **Marco Alberto Javarone**[1]

Executive Summary

The content of this paper is based on a manuscript by the same authors, which will soon be available on arXiv (Hawashin and Jaravone 2025). Quantum computing and artificial intelligence (AI) are often presented as separate technological frontiers. However, it is the synergy of both emerging technologies that makes this one of the most dynamic areas of scientific research and innovation. This work explores the dynamic relationship between these technologies, focusing on how they influence and advance each other.

Background. At its core, quantum computing goes beyond the binary representation of zeros and ones by encoding information in quantum bits, or qubits. A qubit is the quantum analogue of a classical bit; unlike a conventional 0 or 1, it can exist in a superposition of both states simultaneously. This richer representation allows information to be expressed in more complex forms. Two key properties of qubits underpin their computational power: superposition and entanglement. Superposition enables a qubit to embody multiple states at once, while entanglement establishes strong correlations between qubits, even when they are spatially separated. Together, these features allow quantum systems to model and simulate complex phenomena far more efficiently than classical computers. Moreover, the integration of AI into this paradigm enables the development of quantum-inspired approaches. These methods can already deliver practical benefits, even before fully reliable, large-scale quantum machines become available.

Quantum for AI. Encoding information into quantum states introduces fundamentally new paradigms of computation. The distinctive properties of quantum systems—most notably superposition and entanglement—enable quantum computers to process information in ways that surpass the capabilities of classical architectures. In the context of machine learning, this opens alternative approaches to well-known tasks such as classification, regression, clustering, and dimensionality reduction. Although current quantum devices remain small and error-prone, early hybrid schemes that integrate quantum and classical resources have already demonstrated promising results, particularly for data-intensive problems and computationally demanding applications. Ongoing research is testing these methods across diverse domains, including fraud detection in finance, protein-folding prediction in drug discovery, and the optimization of large-scale logistics and supply chains.

AI for Quantum. Quantum computers remain highly fragile devices, as qubits are extremely sensitive to environmental disturbances such as heat, vibrations, or electrical noise. These factors lead to decoherence and errors, making the realization of fully fault-tolerant quantum computation particularly challenging. AI is increasingly being employed to address these limitations. Techniques such as reinforcement learning and neural networks can optimize the design of quantum circuits, automatically calibrate hardware to enhance stability, and even support real-time error detection and correction. Current devices, often referred to as Noisy Intermediate-Scale Quantum (NISQ) machines, typically comprise a few dozen to a few hundred qubits but still exhibit significant error rates. Research in this phase is therefore focused on noise reduction and error mitigation strategies, with the aim of rendering quantum computing more practical and ultimately scalable.

Industry Uses. The integration of quantum computing and AI has already produced early applications across several industries, demonstrating promising potential even in the near term. In finance, hybrid approaches are being explored for fraud detection, portfolio optimization, risk assessment, and market simulation. In healthcare and pharmaceuticals, quantum simulations combined with AI are accelerating drug discovery, improving protein-folding predictions, and enabling the study of complex molecular interactions. In transportation and logistics, quantum algorithms

[1]Exponential Science

are being tested for traffic optimization, battery management, and cargo loading efficiency. In the energy sector, hybrid quantum–classical methods are applied to smart grid management, price forecasting, and large-scale energy simulations. Finally, in telecommunications, efforts are underway to develop quantum-safe encryption and secure communication protocols.

Collectively, these efforts highlight how quantum technologies are beginning to reshape critical industries while laying the groundwork for fault-tolerant and scalable applications, all of which have relied on AI in some capacity to make them practical and effective.

1 Introduction

The emergence of quantum computing has been anticipated for decades, as physicists and theorists reflected on how the principles of quantum mechanics might eventually be translated into computation. The idea of a quantum computer is famously attributed to Richard Feynman, who observed that classical computers would struggle to efficiently simulate quantum systems and argued instead that quantum systems themselves should be used for such tasks (Feynman 1982). Building on this insight, David Deutsch introduced the concept of a quantum Turing machine, offering the first formal description of a universal quantum computer (Deutsch 1985). In essence, this showed that a quantum computer could, at least in principle, simulate any physical system governed by quantum mechanics—for instance, chemical bonding, protein folding, or the behaviour of novel materials at the atomic scale.

For many years, though, the field remained largely theoretical. Classical computing was still rapidly advancing, and artificial intelligence (AI) was gaining momentum, so most of the scientific community paid limited attention to quantum approaches. This changed with Peter Shor's landmark contribution: his algorithm provided the first concrete demonstration, at the theoretical level, that a quantum computer could outperform its classical counterparts (Shor 1997). Since modern cryptographic schemes such as RSA rely on the hardness of integer factorization, Shor's result represented a genuine turning point—often regarded as the work that sparked an entire generation of research.

From that moment, attention increasingly shifted toward building physical devices. Researchers began to explore multiple architectures, including superconducting circuits, trapped ions, photonic systems, and neutral atoms. Each of these approaches carries distinct advantages and limitations, and so far none has established itself as the definitive solution. In many ways, this situation resembles the early days of classical computing, when different designs competed until the transistor emerged as the standard technology.

A fundamental challenge for all current platforms is noise. Qubits are extremely sensitive to their environment, and even minor disturbances can cause errors. Overcoming these issues to achieve a fault-tolerant quantum computer (FTQC) remains one of the hardest problems in the field. For the moment, research is confined to the so-called Noisy Intermediate-Scale Quantum (NISQ) era, where devices typically feature tens to hundreds of qubits but remain error-prone and limited in scope.

This reality has redirected attention toward near-term hybrid approaches, where quantum hardware is combined with classical computing and AI techniques to maximize current capabilities. Hybrid algorithms are already being tested on applications such as the design of new battery materials, protein folding for drug discovery, and large-scale optimization problems. At the same time, AI is proving valuable for the hardware itself, for example in reducing noise, calibrating devices, and improving the reliability of outputs.

The structure of this discussion paper is as follows: Section 2 introduces key quantum concepts and terminology. Section 3 surveys six principal branches of quantum-enhanced machine learning models. Section 4 examines how AI contributes to the development of quantum hardware. Section 5 reviews the different eras of quantum computing, summarizes the main findings, and highlights industry applications. Together, these perspectives illustrate how the intersection of quantum computing and AI is shaping both research and practice. Eventually, Section 6 ends the manuscript.

2 Fundamental of Quantum Technologies

Quantum technologies are often grouped into three main areas: quantum computation, quantum communication, and quantum sensing. While all three are advancing rapidly, the scope of this paper is focused on quantum computation for its relevance to near-term industry application and integration with AI.

As with any other technology, the foundation lies in how information is represented. In classical computing, the basic unit of information is the bit, which can only take one of two values: 0 or 1. In quantum computing, the basic unit is the quantum bit, or qubit. Thanks to the principles of quantum mechanics, a qubit can exist in superposition, meaning it may simultaneously embody aspects of both 0 and 1 until a measurement is performed. This ability to occupy multiple states simultaneously gives quantum computers the potential to explore many possibilities in parallel, rather than sequentially as in classical machines.

Another property that makes quantum computation unique is entanglement. When two qubits become entangled, their states are no longer independent. Instead, they form a partnership in which the measurement outcome of one immediately determines the outcome of the other, even if they are separated by large distances. This phenomenon can be compared to rolling two dice in different rooms and finding that, no matter how many times the dice are rolled, they always show the same number. This analogy is offered only as an aid to intuition: in reality, entanglement is a fundamentally different and much deeper quantum phenomenon, not explainable by simple deterministic correlations. This "spooky action at a distance," as Einstein described it, enables quantum computers to link qubits in ways that unlock powerful forms of processing and communication beyond what is possible in classical systems.

The Bloch sphere is often used to visualize the state of a qubit. A classical bit can only be 0 or 1, like being forced to stand at one of two fixed points, i.e the North Pole or the South Pole of a globe. A qubit, however, is not restricted in this way. Its pure state can be represented as any point on the surface of the sphere, just as a marker can be placed anywhere on Earth: at the equator, in New York, in Tokyo, or at the poles. The North Pole corresponds to state 0, the South Pole to state 1, but the entire surface in between the poles represents all the possible superpositions of these two states.

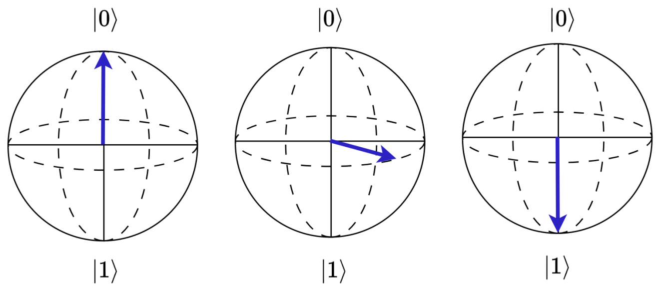

Figure 1: Representation of three different quantum states as vectors on the Bloch sphere.

The promise of quantum computing is matched by its difficulty. Qubits can lose their quantum state on extremely short timescales when exposed to even minor disturbances from their surroundings. This problem, known as decoherence, makes it extremely challenging to keep information stable long enough to perform useful calculations. In classical computing, errors are easier to manage: information can be duplicated, and redundancy can be used to detect and correct mistakes. In the quantum world, however, this approach is impossible because of a rule in physics called the no-cloning theorem. Attempting to copy a quantum state directly would disturb it, causing the state to collapse and the original information to be lost.

To overcome this, scientists have developed a method called quantum error correction. Instead of storing information in a single qubit, it is redundantly encoded across multiple entangled qubits.

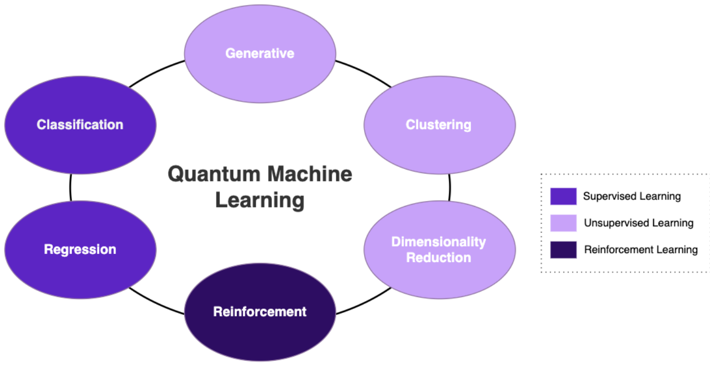

Figure 2: Diagrammatic representation of the six quantum machine learning paradigms identified and investigated by this paper.

If one of them is disturbed, the others can be used to detect and correct the error, preserving the quantum information without collapsing it.

Today's NISQ devices contain tens to hundreds of qubits and can already perform limited computations, but they remain prone to noise and error. While these machines are not yet capable of delivering large-scale quantum advantage, they are invaluable for experimentation. They allow researchers to test hybrid algorithms, in which classical optimization is paired with quantum subroutines, and to explore early applications in fields such as material science, drug discovery, optimization, and many more.

3 Quantum for AI Software

Recent advancements have integrated quantum technologies with AI, particularly quantum machine learning (ML) paradigms. Here, this mirrors the classical workflow of data being prepared, processed, and output. Specifically, quantum embedding is used to map classical data to quantum states, manipulation of gate sequence for processing, and the qubit is measured for the output.

One of the biggest advantages of quantum computing comes from the way it represents information in a Hilbert space. The Bloch sphere illustrates the state space of a single qubit. This space allows quantum computers to handle far larger amounts of information with far fewer resources than classical computers. The difference grows dramatically as the problem gets bigger. With just 5 qubits, a quantum computer can represent $2^5 = 32$ different dimensions. With 10 qubits, it can represent $2^{10} = 1024$ dimensions. And with 20 qubits, the system can represent over one million dimensions at once. This exponential scaling is what makes quantum computing so powerful: relatively small quantum systems can represent spaces that would overwhelm even the largest classical high-performance computers.

Quantum-enhanced machine learning applications are a maturing field of research. Significant progress has even allowed several concepts to move past theoretical proof-of-concept stages to the development of practical applications. This section maps familiar classical tasks (i.e. classification, regression, clustering, ..) to quantum-appropriate solutions and discusses ways in which this hybrid approach makes it more applicable to industry applications.

Access to a quantum computer remains beyond the reach of most individuals today. Even when available, current quantum devices are often too noisy to reliably perform computations. A practical approach for the use of near-term quantum devices is through hybrid quantum-classical mechanisms. Parametrized Quantum Circuits (PQCs) have been employed in numerous studies to demonstrate true quantum advantage.

Supervised learning refers to the task of training a model on pairs of inputs and outputs, allowing it to generalize to unseen data. The two most common tasks within this paradigm are classification and regression, both of which have natural extensions into quantum machine learning.

Classification is the task of sorting inputs into distinct categories - for example, deciding whether an email is spam or not. Classical methods, such as support vector machines or neural networks, achieve this by learning a boundary that separates one class from another based on shared features. Quantum computing approaches the same problem in a different way. Instead of working within the limits of classical space, quantum systems can place data into a much larger mathematical space, the Hilbert space, where patterns and separations that are hidden to classical methods can become visible. One way this is done is through a quantum version of the classical kernel method, which measures the similarity between data points in this expanded space. Another approach uses variational quantum classifiers (Zhou et al. 2023) (Miyahara and Roychowdhury 2022). Here, the data is encoded into a quantum circuit by rotating and entangling qubits, and the results of measurements are interpreted as class labels. The model is trained much like a classical one: a cost function is used to evaluate performance, and a classical optimizer iteratively updates the quantum circuit's parameters until the model can reliably distinguish between categories.

Regression, unlike classification, is about predicting continuous values rather than assigning data to fixed categories. A familiar example is predicting house prices based on features like size, location, and number of rooms. Classical regression methods work by finding a function that best fits the data while minimizing the difference between predictions and actual outcomes. Quantum regression follows a similar principle but uses the unique behaviour of quantum states to represent and process very complex relationships in data. One promising approach, known as quantum linear regression, draws on quantum routines that can solve certain mathematical problems much faster than classical computers, at least in theory. This could allow quantum systems to handle very large datasets and high-dimensional problems far more efficiently. In practice, today's quantum hardware is still noisy, so researchers rely on hybrid approaches (PQCs) with adjustable parameters that are tuned during training (Suzuki and Katouda 2020).

Unsupervised learning is a way for computers to find patterns in data without being told what the "right answer" is. Instead of working with labelled examples, the goal is to uncover the hidden structure within the data. One of the most common tasks here is clustering, where similar items are grouped together based on how close or related they are.

Clustering approaches to quantum computing include the q-means methods, the quantum version of the well-known k-means algorithm DiAdamo et al. 2022. In this method, data is grouped into clusters by calculating the "centre" (centroid) of each group. Quantum computers can estimate similarity, or fidelity, between quantum states using the SWAP test. Once the similarity of each item to each cluster is measured, the data points are reassigned to their closest cluster, and the centroids are updated. This process repeats until the clusters settle into stable groups.

Dimensionality reduction tackles the "curse of dimensionality", where data becomes exponentially more complex as the number of features increases, quickly overwhelming classical methods. Quantum systems can handle this more efficiently by mapping data into a Hilbert space, though further optimization is needed. A breakthrough came from Lloyd et al, who introduced quantum Principal Component Analysis (PCA), a method akin to summarizing a long book into its key chapters (Lloyd et al. 2014). They showed that a quantum computer could potentially perform this summarization faster under certain assumptions about data access. Later work noted that the original method became too costly for complex datasets, leading to an improved approach that allows quantum computers to adjust key values directly without computing every detail, making it more practical at scale (Nghiem 2025). Other directions include quantum auto encoders, which compress data by keeping only essential features (Romero et al. 2017). A useful analogy is reducing a high-resolution photo to a smaller file without losing its main content. In practice, this is achieved with PQCs trained to capture the most relevant information while filtering out noise.

Generative Modelling in classical machine learning has inspired similar efforts in quantum computing. Quantum Born Machines create samples directly from the natural probabilities of quantum states, making them relatively straightforward to run on today's devices. Quantum Boltzmann Machines, on the other hand, try to model data using energy landscapes, which requires more complex quantum states that today's hardware cannot yet handle well. Another key approach is Quantum Generative Adversarial Networks (Q-GANs), which, like their classical counterparts, pit a generator against a discriminator to learn data distributions. In quantum setups, the generator can be a Born Machine or a PQC, while the discriminator may be classical or quantum. Quantum Variational Autoencoders (Q-VAEs) take a different path, compressing data into simpler forms and reconstructing it to generate new samples. Depending on the task, the encoder and decoder can be quantum, classical, or a mix of both. Early studies show Q-VAEs can outperform classical models on benchmarks like MNIST (Khoshaman et al. 2018). Finally, quantum transformers are gaining traction, especially in language processing. Words are encoded as quantum states, with self-attention handled through quantum operations and positional encoding added either classically or within circuits. Recent work includes hybrid vision transformers and fully quantum-native models such as Quixer, with initial results suggesting they already match classical baselines (Khatri et al. 2024).

Reinforcement learning is a method where an agent learns by interacting with its environment: taking actions, receiving rewards, and improving decisions over time. In quantum RL, early work on model-based methods, which require building a full model of the environment, has proven too demanding for current hardware. Most research now focuses on model-free methods. In value-based RL, quantum circuits are used to estimate the long-term rewards of actions, while in policy-based RL, they directly learn strategies by producing and refining probability distributions over actions. A combined approach, the actor–critic method, uses a quantum "actor" to propose actions and a classical or quantum "critic" to evaluate them. Recent studies have even applied this hybrid model to real-world tasks such as securing power grids, demonstrating early steps of quantum RL beyond theory (Peter and Korkali 2025).

4 AI for Quantum Hardware

The development of quantum computers faces numerous challenges. Classical simulations have enabled early demonstrations of quantum computation, but there is a broad consensus that achieving FTQC is nearly impossible without advanced optimization techniques. AI offers powerful tools for error correction, error mitigation, and quantum system design, all of which contribute to advancing the current state of quantum computers.

Circuit Complexity plays a central role in determining the performance of a quantum algorithm. It is defined by the minimum number of quantum gates required to implement a given unitary transformation. The total number of gates is referred to as the circuit size, while the number of sequential layers of gates is known as the circuit depth. Minimizing both size and depth is crucial for making quantum computing more practical. Artificial intelligence has become an important tool in this work. Generative AI models (normally used for language tasks) can be trained on datasets of circuits to propose efficient new designs. Recent studies have used transformer- and diffusion-based models to generate low-energy or hardware-efficient circuits, and variational autoencoders to identify effective circuit structures (ansätze) for specific problems (Nakaji et al. 2024). RL has also been applied to circuit design. Here, an agent builds circuits step by step, adding or adjusting gates based on a reward signal such as accuracy or efficiency. RL has been used for parameter transfer, where instead of starting circuits with random parameters, they are initialized using values learned from similar problems, therefore avoiding wasted effort and speeding up convergence (Verdon et al. 2019). Together, these approaches highlight how combining AI with quantum theory can make circuit design more efficient, scalable, and tailored to both hardware limits and practical applications.

Error Detection and Correction presents a fundamental problem with realizing a FTQC. Research on handling quantum errors and thereby enabling reliable quantum computation remained largely theoretical until the development of quantum error correction codes. This is based on the concept of encoding one logical qubit into many physical qubits, thereby protecting information against errors (Shor 1995). In order to correct an error, it must first be accurately identified. This can be categorized into three main types of quantum errors: bit-flip, phase-flip, and simultaneous

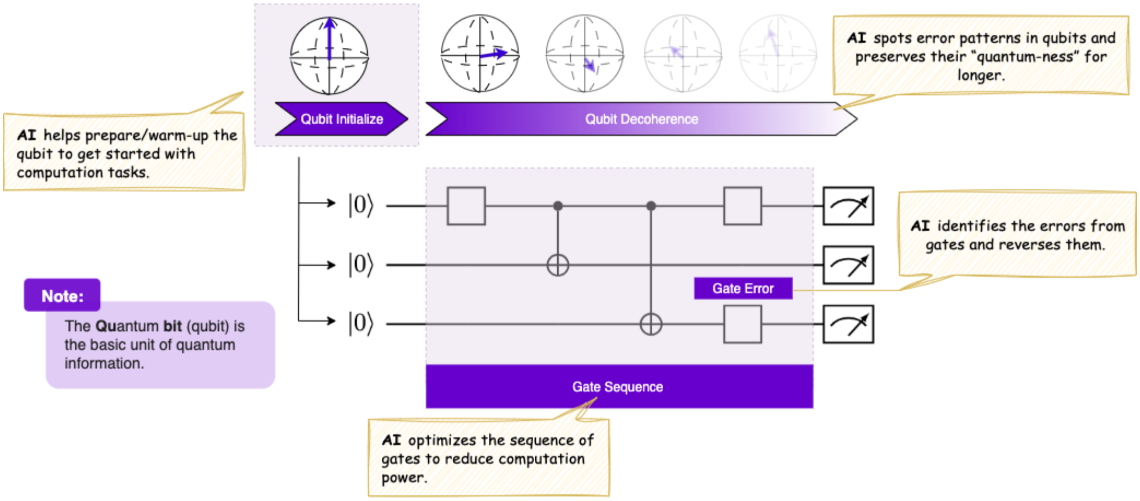

Figure 3: AI-assisted quantum computing pipeline. AI supports qubit initialization, mitigates decoherence, detects and corrects gate errors, and optimizes gate sequences to preserve quantum states and enhance computational efficiency.

bit and phase errors. There are several quantum error correction codes, the simplest of which is repetition codes inspired by classical error correction through storing multiple redundant qubits. However, this only corrects bit-flip errors and is inefficient to scale because it requires many physical qubits. Stabilizer codes represent a broad family of quantum error correction codes (e.g., Shor's code, Steane code, Surface codes, etc.), all of which share the same underlying framework. Instead of simple copies, these codes arrange groups of qubits in structured patterns that can automatically detect and correct a wider range of errors.

5 Industry Applications

The majority of industries are already well aware of the advantages and risks associated with the eventual realization of a fully FTQC. In the current state of the world, this is reflected by collaborations between service providers and quantum computing companies. Because of this, there exists a plethora of both press releases and corresponding papers published on the research being done. The following sections are grouped by industry and highlight how large organizations within these sectors have adapted and begun implementing quantum technologies into their workforce.

In **finance**, computational time and accuracy often directly impact profit-and-loss margins. Leading financial institutions quickly recognized this opportunity and have made significant investments in quantum technologies. Portfolio optimization has been a key focus: BBVA partnered with Multiverse Computing to test D-Wave's quantum annealer for investment strategies, while IQM and DATEV experimented with a 20-qubit gate-based system running hybrid optimization-based algorithms to improve product portfolios. Goldman Sachs and IonQ explored quantum-enhanced Monte Carlo simulations for derivative pricing, while JPMorgan Chase used Quantinuum's random circuit sampling to generate certified randomness for stochastic modeling. Mastercard has also entered the space, testing hybrid quantum-classical methods with D-Wave to refine algorithmic trading and market prediction strategies.

Healthcare and pharmaceutical companies face heavy computational demands in simulating molecules and discovering new drugs. Moderna partnered with IBM with the aim of applying quantum and AI methods to predict molecular properties, improving the design of mRNA therapies. Biogen, Accenture, and 1QBit built a quantum-powered comparison tool to study molecular differences linked to neurological diseases like Alzheimer's and Parkinson's. AstraZeneca pushed this further by combining IonQ's quantum processor with NVIDIA GPUs on AWS, reporting significant acceleration of chemical reaction simulations compared to classical methods.

Transportation and logistics companies have used quantum to tackle large optimization problems. Volkswagen and D-Wave famously demonstrated real-time bus routing in Lisbon, proving that a quantum annealer could adapt to live traffic data. Hyundai turned to IonQ to simulate lithium compounds for better EV battery design, while Airbus applied a quantum algorithm to cargo loading, finding efficient ways to distribute weight and volume within an aircraft. These projects showed how quantum computing can handle problems that quickly overwhelm classical solutions.

In the **energy sector**, providers have begun to explore quantum methods for pricing and grid management. E.ON, one of Germany's largest utilities, worked with IBM to model complex consumer pricing schemes and later joined the Q-GRID project to test grid optimization using quantum annealing. In Italy, Eni collaborated with PASQAL to add neutral-atom quantum processors into its high-performance computing workflows, with the goal of accelerating large-scale energy simulations through hybrid quantum–classical approaches.

Telecommunications companies, while not yet deeply integrating AI into their quantum efforts, have moved early to secure their networks. Nokia partnered with Turkcell to deploy quantum-safe IPsec encryption, while e&, the UAE's major telecom provider, rolled out a quantum-safe network architecture. Deutsche Telekom has taken a more experimental route, working with Qunnect on research toward building a quantum internet.

6 Conclusions

In the NISQ era, with its limited hardware capabilities, AI has allowed for advanced uses of quantum for near-term practical applications. In many ways, this has improved both software and hardware developments. In this paper, the authors identify what encompasses the relation between Quantum and AI, specifically along two main directions: Quantum for AI software, AI for quantum hardware. With the purpose of providing a holistic explanation of quantum machine learning paradigms, the paper identifies 6 key tasks: classification, regression, clustering, reinforcement learning, generative learning, and dimensionality reduction. These tasks utilise the hybrid architecture of variational algorithms that leverage quantum principles (superposition, entanglement and interference), and are optimized via classical cost functions and optimizers. The need for this classical-quantum approach is due to the difficulty of deploying quantum-native algorithms to the current state of quantum hardware. This integration has, in turn, positively progressed the application of quantum methods to initial use cases within industries such as finance, healthcare, transportation, etc. As for quantum hardware, their sensitivity to the external environment presents new challenges leading to noise, decoherence, and operational errors (i.e phase-flip, bit-flip) that limit scalability. The integration of classical AI methods, such as reinforcement learning, generative modelling, and neural-network-based decoders, has allowed for significant progress to be made, primarily in error detection and circuit complexity optimization. The roadmap towards a FTQC is advancing steadily, and it is with the reliance on other advanced technologies such as AI to exploit the boundaries of the resources given today to accelerate tomorrow's technology.

References

Deutsch, David (1985). "Quantum theory, the Church–Turing principle and the universal quantum computer". In: *Proceedings of the Royal Society of London. A. Mathematical and Physical Sciences* 400, pp. 97–117. DOI: 10.1098/rspa.1985.0070. URL: https://doi.org/10.1098/rspa.1985.0070.

DiAdamo, Stephen, Corey O'Meara, Giorgio Cortiana, and Juan Bernabe-Moreno (2022). "Practical Quantum K-Means Clustering: Performance Analysis and Applications in Energy Grid Classification". In: *IEEE Transactions on Quantum Engineering* 3, pp. 1–16. ISSN: 2689-1808. DOI: 10.1109/tqe.2022.3185505. URL: http://dx.doi.org/10.1109/TQE.2022.3185505.

Feynman, Richard P. (1982). "Simulating physics with computers". In: *International Journal of Theoretical Physics* 21.6-7, pp. 467–488. DOI: 10.1007/BF02650179.

Hawashin, Hala and Marco Jaravone (2025). "A Classical Perspective on the Synergies of Quantum and Artificial Intelligence". N.p.: arXiv.

Khatri, Nikhil, Gabriel Matos, Luuk Coopmans, and Stephen Clark (2024). *Quixer: A Quantum Transformer Model.* arXiv: 2406.04305 [quant-ph]. URL: https://arxiv.org/abs/2406.04305.

Khoshaman, Amir, Walter Vinci, Brandon Denis, Evgeny Andriyash, Hossein Sadeghi, and Mohammad H Amin (Sept. 2018). "Quantum variational autoencoder". In: *Quantum Science and Technology* 4.1, p. 014001. ISSN: 2058-9565. DOI: 10.1088/2058-9565/aada1f. URL: http://dx.doi.org/10.1088/2058-9565/aada1f.

Lloyd, Seth, Masoud Mohseni, and Patrick Rebentrost (July 2014). "Quantum principal component analysis". In: *Nature Physics* 10.9, pp. 631–633. ISSN: 1745-2481. DOI: 10.1038/nphys3029. URL: http://dx.doi.org/10.1038/nphys3029.

Miyahara, Hidetoshi and Vinay Roychowdhury (2022). "Ansatz-Independent Variational Quantum Classifiers and the Price of Ansatz". In: *Scientific Reports* 12. Received 19 September 2021, Accepted 16 September 2022, Published 14 November 2022, p. 19520. DOI: 10.1038/s41598-022-20688-5. URL: https://doi.org/10.1038/s41598-022-20688-5.

Nakaji, Kouhei, Lasse Bjørn Kristensen, Jorge A. Campos-Gonzalez-Angulo, Mohammad Ghazi Vakili, Haozhe Huang, Mohsen Bagherimehrab, Christoph Gorgulla, FuTe Wong, Alex McCaskey, Jin-Sung Kim, Thien Nguyen, Pooja Rao, and Alan Aspuru-Guzik (2024). *The generative quantum eigensolver (GQE) and its application for ground state search.* arXiv: 2401.09253 [quant-ph]. URL: https://arxiv.org/abs/2401.09253.

Nghiem, Nhat A. (2025). *New Quantum Algorithm for Principal Component Analysis.* arXiv: 2501.07891 [quant-ph]. URL: https://arxiv.org/abs/2501.07891.

Peter, Benjamin M. and Mert Korkali (2025). *Quantum-Enhanced Reinforcement Learning for Power Grid Security Assessment.* arXiv: 2504.14412 [eess.SY]. URL: https://arxiv.org/abs/2504.14412.

Romero, Jonathan, Jonathan P Olson, and Alan Aspuru-Guzik (Aug. 2017). "Quantum autoencoders for efficient compression of quantum data". In: *Quantum Science and Technology* 2.4, p. 045001. ISSN: 2058-9565. DOI: 10.1088/2058-9565/aa8072. URL: http://dx.doi.org/10.1088/2058-9565/aa8072.

Shor, Peter W. (Oct. 1995). "Scheme for reducing decoherence in quantum computer memory". In: *Phys. Rev. A* 52 (4), R2493–R2496. DOI: 10.1103/PhysRevA.52.R2493. URL: https://link.aps.org/doi/10.1103/PhysRevA.52.R2493.

— (Oct. 1997). "Polynomial-Time Algorithms for Prime Factorization and Discrete Logarithms on a Quantum Computer". In: *SIAM Journal on Computing* 26.5, pp. 1484–1509. ISSN: 1095-7111. DOI: 10.1137/s0097539795293172. URL: http://dx.doi.org/10.1137/S0097539795293172.

Suzuki, Teppei and Michio Katouda (2020). "Predicting toxicity by quantum machine learning". In: *Journal of Physics Communications* 4.12, p. 125012. DOI: 10.1088/2399-6528/abd3d8. URL: https://iopscience.iop.org/article/10.1088/2399-6528/abd3d8.

Verdon, Guillaume, Michael Broughton, Jarrod R. McClean, Kevin J. Sung, Ryan Babbush, Zhang Jiang, Hartmut Neven, and Masoud Mohseni (2019). *Learning to learn with quantum neural networks via classical neural networks.* arXiv: 1907.05415 [quant-ph]. URL: https://arxiv.org/abs/1907.05415.

Zhou, J., D. Li, Y. Tan, et al. (2023). "A multi-classification classifier based on variational quantum computation". In: *Quantum Inf Process* 22. Received 14 April 2023, Accepted 06 October 2023, Published 17 November 2023, p. 412. DOI: 10.1007/s11128-023-04151-6. URL: https://doi.org/10.1007/s11128-023-04151-6.

Exponential Science

Exponential Science integrates exponential technologies to address complex global challenges through education, research, and innovation. Led by Dr. Paolo Tasca and Nikhil Vadgama, Exponential Science collaborates with academia, industry, and government to foster deep tech knowledge, generate pioneering research, and promote early-stage innovations. Exponential Science, as a natural evolution of the DSF's work, aims to harness the convergence of technologies such as blockchain, AI, and IoT to create transformative impacts towards a more decentralized and open society.