Quantum + AI = Quantum AI

Maria Spiropulu & Hartmut Neven

Advanced AIs of the future will harness the most powerful computational operations known today: quantum computations. The synthesis of quantum computing and AI will give rise to Quantum AI – artificial intelligence with access to quantum computational resources – which is poised to expand the capabilities of intelligent systems. This advancement will lead to transformative abilities, such as ultraprecise sensing, the realistic simulation of complex natural phenomena, the efficient solution of classically intractable mathematical problems, or protocols for secure communication and enhanced coordination in multi-agent systems. The impact of this synthesis will be profound, promising to accelerate scientific discovery in areas from molecular biology to quantum gravity, while empowering engineers to create novel pharmaceuticals, viable fusion reactors, or designer materials with desirable properties. Quantum AI might also help to address a critical bottleneck in classical machine learning: its escalating demand for data and computational power. Beyond such applications lies another even more consequential possibility. If, as some recent theories and data from quantum neurobiology suggest, quantum processes give rise to consciousness and free will, then only Quantum AI – unconstrained by the limitations of classical computation – could enable the expansion of human consciousness and a transcendence of our biological origins.

Quantum computing and artificial intelligence are often described as the most transformative computational technologies we will see develop during our lifetimes. The recent 2025 Nobel Prize in Physics (awarded to John Clarke, Michel Devoret, and John M. Martinis for work on macroscopic quantum effects), alongside the 2024 Nobel Prizes in Physics (awarded to Geoffrey Hinton and John Hopfield for their contributions to the study of artificial neural networks) and Chemistry (awarded to Demis Hassabis, John Jumper, and David Baker for solving the protein-prediction problem and creating novel proteins using AI systems), recognized the profound breakthroughs in both quantum technologies and the neural networks that power AI systems. However, today they remain largely distinct technologies. Artificial intelligence takes in data, creates models, and utilizes those models to achieve objectives. Quantum computing aims to develop a new type of hardware that harnesses the potent computational operations allowed by quantum physics – such as entanglement and superposition – replacing the Boolean

https://doi.org/10.1162/DAED.a.985

(binary) logic implemented on today's semiconductor-based hardware. While AI is already a useful tool to design and build quantum computers, less appreciated is the fact that AI based on hardware with access to quantum resources will have capabilities inaccessible to AI running solely on classical hardware. The thesis we put forth in this essay is that quantum computing and AI will merge to become Quantum AI, enabling discoveries across science and a range of breakthrough technical capabilities. By leveraging quantum advantages, artificial intelligence will utilize some of the most potent computational operations nature affords us to improve our understanding of the universe and natural intelligence. Here, we envision the impact of synthesizing these breakthroughs.

Artificial intelligence as practiced today is nearly synonymous with machine learning, relying on the ability to ingest huge amounts of data to create powerful models. However, many foundational computational tasks are challenging for classical AI to learn due to the impracticality of obtaining sufficient training data and doing inference on it. Consider factoring, an algebraic task relevant to cryptography. Empirical evidence – about fifty years of failure to find an efficient classical factoring algorithm – suggests that AI unaided by quantum computers will not be able to solve this task efficiently. The difficulty is a scaling challenge: the training cycles and inference compute required grow exponentially with the number of bits needed to represent the integer being factored. To calibrate the reader, factoring a 2,048-bit number would take trillions of years on a classical supercomputer but would only take hours on a quantum computer.

Classical AI also struggles to solve hard combinatorial optimization problems, which are ubiquitous in engineering, finance, and computer science. Again, empirical evidence shows that finding the optimal solution to such problems takes time that grows exponentially with the size of the problem. Sometimes even finding a good approximation within a reasonable time span can be impractical, a situation mathematicians call "APX hard."[1] Machine-learning algorithms struggle in this domain because finding the global optimum is a "needle in a haystack" problem, requiring the identification of a specific solution within a vast combinatorial search space. However, theoretical analysis has shown that quantum computers within a given amount of time can find higher-quality solutions than their classical counterparts. For many problem classes, a quantum computer can reach the same quality approximation in quadratically fewer steps, and sometimes in exponentially fewer steps, as the recently published Decoded Quantum Interferometry (DQI) algorithm and new results for the Quantum Approximate Optimization Algorithm (QAOA) show.[2] If we compare two AI systems but one of them has access to a quantum machine giving us high-quality solutions to optimization that are classically inaccessible, then the quantum-enhanced system would be more valuable. Today, there are at least seventy known algorithms for which quantum computers offer algorithmic scaling advantages, many of them useful for AI, and we expect many more will be discovered.[3] A recent XPRIZE challenge that called for submissions of useful beyond-classical algorithms yielded three hundred and twenty entries.[4]

Complexity theorists attempting to map out the sets of problems that can be efficiently solved by quantum and classical algorithms draw the relationship as shown in the diagram in Figure 1. In complexity theory, a problem is considered solvable when the resources required, such as the number of algorithmic steps or the amount of memory, only grow polynomially with the size of the problem. But as is often the case in complexity theory, while considered highly probable, we still lack complete proof that all these relationships are correct.

Quantum advantage resides in the problem set that is quantumly solvable but classically intractable: the regions of BQP (quantumly solvable) and BQP/qsamp (quantumly solvable with quantum data) that do not overlap with BPP (classically solvable) or P/samp (classically solvable with classical data). A prime example of this divergence is found in out-of-time-order correlators (OTOCs), probes used to characterize how information spreads in complex quantum systems. While OTOC computation scales exponentially on classical hardware, it was recently demonstrated that the Willow quantum processor can perform this calculation efficiently.[5] This computation is useful for learning the structure of molecules and solids from spin spectroscopy data, as obtained by nuclear or electron magnetic resonance, and constitutes the first known example of an AI-relevant algorithm for which a quantum computer outperformed classical supercomputers.

As a prime example of how data from quantum computers can empower classical AI, let's look at the accurate modeling of complex physical systems. This is known to have applications ranging from the discovery of novel materials and pharmaceuticals to a better fundamental understanding of our universe. AI's main quest, constructing models from data it ingests, is impeded by the fact that nature is governed by the laws of quantum physics. Integrating the equations of quantum mechanics typically comes at an exponential cost for classical computers, making them an ill-suited tool to model nature at the most fundamental level. In contrast, quantum computing's ability to simulate nature from the ground up gives access to a broader, more detailed set of training data for AI. Consider computational chemistry, where popular methods involve machine learning of atomistic force fields or density functionals using data from *ab initio* electronic structure calculations.[6] Such approaches enable fast, accurate calculations of complex molecular structures from manageable amounts of training data. However, the model's output accuracy is only as good as the training data's accuracy. When generated classically, the data's quality will be limited by the challenge of modeling strongly correlated electrons – an inherently quantum-mechanical problem. Therefore, data derived from high-accuracy quantum calculations represent a valuable resource for training advanced chemistry models.

Figure 1
Problems That Can Be Efficiently Solved by Quantum Algorithms versus Classical Algorithms

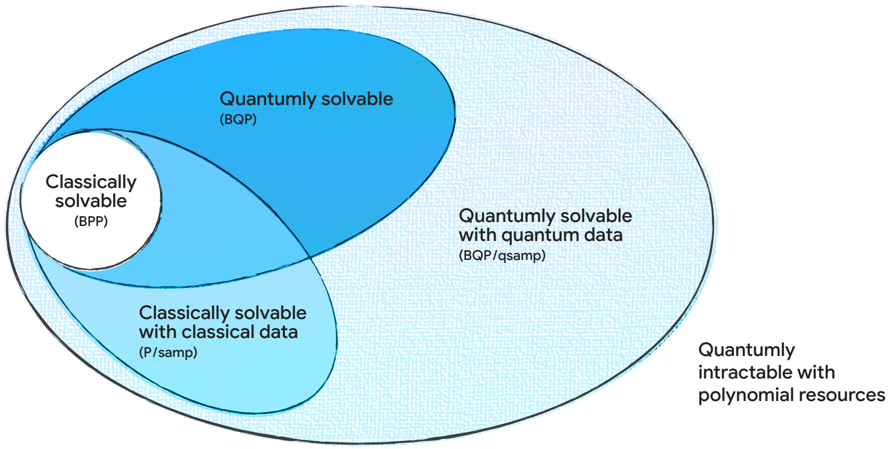

Quantum computers are expected to solve a larger set of problems than classical computers. Access to data, or *advice* as complexity theorists call it, makes classical and quantum algorithms more powerful. However, there remain problems that even a quantum computer with access to data will not be able to solve efficiently. The complexity theory acronyms in Figure 1 have the following meanings: BPP = Bounded-error Probabilistic Polynomial time. P/samp = Polynomial time with classical data. This is the set of all problems solvable on classical computers; it is also a superset of BPP because data make algorithms more powerful. BQP = Bounded-error Quantum Polynomial time. BQP/qsamp = Bounded-error Quantum Polynomial time with quantum data. This is the set of all problems solvable on quantum computers. It is expected to be strictly larger than P/samp. Again, it is a superset of BQP. Source: Figure by the authors. On access to data in quantum algorithms, see Hsin-Yuan Huang, Michael Broughton, Masoud Mohseni, et al., "Power of Data in Quantum Machine Learning," *Nature Communications* 12 (1) (2021): 2631.

Classical AI has in recent years found great success in predicting how proteins fold.[7] But the challenge remains when the amount or quality of training data is insufficient. Consider the case of a protein bound to a ligand or a membrane, both common scenarios in cell biology. The data scarcity occurs partly because magnetic resonance spectroscopy, a primary tool for gathering molecular structure information, often fails in these situations. Other tools are available, such as X-ray crystallography or cryogenic electron microscopy, but even if applicable, they must be painstakingly adapted for different proteins or cells. In general, access to large, high-quality datasets is the exception rather than the norm in the life sciences. It is noteworthy that the critical dataset for the protein-folding problem, the Protein

Data Bank, has been assembled over fifty years. With more varied and larger sets of training data, the performance of prediction tools such as AlphaFold could be improved. Quantum computers can help generate such training sets. Nuclear spins in a folded protein generate complex quantum dynamics, even at room temperature and even when the protein's electronic structure does not have high entanglement. A recent collaboration involving the Google Quantum AI team and researchers at UC Berkeley demonstrated that these spin dynamics can reveal geometric details of small organic molecules.[8] Scaling this approach to proteins could provide the necessary distance data to determine accurate three-dimensional protein structures. Such analysis, which involves the computation of OTOCs, is feasible for a quantum computer capable of interpreting spectroscopic data from spin configurations but is likely to remain out of reach for classical biochemistry methods.

The time evolution of quantum systems is another canonical example of data that quantum computers can efficiently produce but that are typically intractable to simulate classically. There are also many contexts in which quantum dynamics are extraordinarily difficult to probe experimentally, thus starving classical learning models of the data they would need to model quantum dynamics, even presupposing that classical models can learn those dynamics at all. Consider the dynamics of an electron plasma in the warm dense matter regime as it slows down a high-energy alpha particle flying through it. Understanding this drag as a function of temperature is critical for correctly modeling the preignition stage of inertial confinement fusion reactors, such as those at the National Ignition Facility that have achieved scientific net energy gain. Unfortunately, getting precise experimental data on this involves temperatures and pressures comparable to those in Jupiter's core. Such experiments only last nano- to picoseconds and, given the costs of prototype fusion reactors, such extreme conditions would be extraordinarily expensive to replicate: hundreds of trillions of dollars for a one-second stream of data. As a consequence, one cannot even accurately measure the temperature of the fuel at that stage, let alone probe the electron dynamics. But simulations of these conditions have been shown to be feasible for quantum computers.[9] If one ever hopes to build machine-learning models of matter in such extreme conditions, one will almost certainly require a quantum computer to accurately generate the data and to attempt to probe them experimentally. Even in regimes that are much more accessible experimentally, such as those in room temperature chemistry, quantum simulation enables much more precise and nuanced virtual measurements on the physical system than would ever be possible in a laboratory.

In the context of generating training data, one can think of a quantum computer as a physics or chemistry lab that can be programmed to obtain data from experiments that are difficult or even infeasible to execute by any other means. Interestingly, quantum computers can sometimes offer advantages even for modeling systems that are described effectively by classical physics. Examples include newer algorithms promising exponential speedups over classical methods for solving certain types of differential equations and specific linear algebra problems that could enable the simulation of molecular mechanics, the structural health of buildings, or wave propagation in solids.[10]

Physics simulations on a quantum computer can be utilized in tandem with AI to extend investigations into quantum gravity. Here, entanglement – the sharing of quantum information between systems regardless of spatial separation – is the key resource for rephrasing questions of space-time. This connection is formalized in the *ER* = *EPR* hypothesis proposed by theoretical physicists Juan Maldacena and Leonard Susskind in 2013, which posits that quantum entanglement and the geometry of space-time are deeply linked.[11] A collaboration between Caltech, Harvard, Fermilab, and Google Quantum AI experimentally demonstrated this link using traversable wormholes, which can occur in semiclassical quantum gravity and exhibit features related to both holography and quantum teleportation.[12] The project utilized a quantum circuit implemented on a Google Sycamore processor to simulate a simplified, dual description of a traversable wormhole, successfully observing the teleportation dynamics predicted by semiclassical gravity. Features of wormhole dynamics observed on the quantum processor can then be used to train AI models. These models, guided by a loss function favoring simpler descriptions of gravity, can then predict a larger family of physical models.[13] These AI-generated predictions can subsequently be realized in detail on quantum hardware, serving as an experimental testing ground for holography and wormhole physics that lies beyond the scope of analytically tractable methods.

Physicists argue that because the objects we observe in nature are generally governed by effective, classical laws, these objects can be learned by classical AI. This is only partially correct. For example, certain static properties of a molecule in its lowest energy state – called the ground state – can sometimes, although not always, be approximated well with classical methods such as tensor network techniques.[14] But when it comes to dynamics, classical methods alone are insufficient: the precision of quantum mechanical calculations is necessary.[15] Any system created by nature was almost certainly created efficiently. Since nature is quantum in its core and origins, simulating that process on a quantum computer must also be efficient. The examples that classical AI can handle are only a subset of systems in nature that science has considered so far. As science investigates new and not yet well-understood systems – such as dark energy, dark matter, black holes, and their connection with the origin of space and time and gravity – a quantum computer will almost certainly be necessary. Some of the future objects of study may not be found in nature at scales accessible with current technologies but can nonetheless be efficiently created in a lab with the guidance of Quantum AI, such as high temperature superconductors and other "designer" states of matter. The small wormhole model discussed above can be considered a Quantum AI–proposed "new state of matter."

Sensing and processing quantum data reveal previously invisible facets of our universe. Without the ability to interact directly with quantum superposition states, classical AI can handle only low-dimensional projections, obscuring the intricate high-dimensional structure such states may contain. However, quantum computers open the possibility of sensing and processing quantum data directly, circumventing the loss of information that occurs when quantum states are measured classically. In a 2022 paper published in *Science*, the Google Quantum AI team with researchers from Caltech, Harvard, and other research organizations showed that the ability to coherently process quantum states on a quantum computer offers the capability to learn structure in quantum data with exponentially fewer training examples compared with classical methods – an information-theoretic gap that cannot be overcome.[16] Combine this with any kind of time limitation on sample collection, and this phenomenon quite literally renders aspects of the physical world invisible without access to a quantum computer. Moreover, this was shown experimentally on the still-noisy quantum processors we have today, meaning this type of processing could be integrated with quantum sensors in the future without many algorithmic changes. To make this a practical reality, quantum sensing that uses entangled probes as well as transduction of the quantum information to the processor are still required. Google recently funded a program in quantum transduction to bring this exciting technology pathway to fruition.[17] Recent work has shown that quantum processing of even classical signals is required to achieve optimal detection.[18]

Historically, until around the early 2000s, scientific observations were exclusively recorded as classical data, such as hand-drawn sketches or photographic images from telescopes and microscopes. Quantum sensing offers an alternative paradigm: incoming particles or fields directly alter the quantum state of a set of qubits (short for quantum bit, the fundamental unit of information in quantum computing).[19] This quantum information is then processed by a quantum computer, a setup that can convey exponential advantages: for instance, to learn properties of the system under observation from exponentially fewer training examples than required by a learner that ingests classical data.

The ability to process quantum-sensing data directly could lead to extraordinary advances in many areas of experimental science. For example, laboratory detection of axion dark matter may require sensor arrays with noise reduction orders of magnitude better than the standard quantum limit, the fundamental limit of precision in a continuous quantum measurement when using classical resources. Prototype systems have demonstrated this via quantum nondemolition readout of the sensors using superconducting qubits, which in turn would allow additional processing with quantum circuits.[20] Another example is quantum sensor arrays deployed at a future particle collider, such as the Future Circular Collider under consideration at CERN. The combined requirements of high intensity for the colliding beams and extraordinary precision in determining position and timing of produced particles will be made even more challenging when searching for exotic new particles that interact very weakly with detector components. Quantum sensor arrays embedded in larger collider detectors offer uniquely low energy-detection thresholds, nanoscale position resolution, and picosecond timing resolution.[21] Such future particle colliders will serve as a scaled-up platform for discovery by generating massive amounts of quantum-native data that can only be fully distilled through end-to-end AI and quantum workflows. By utilizing quantum sensor arrays to capture entanglement data, they will provide the high-fidelity training sets necessary for Quantum AI workflows to construct the ultimate digital twin of nature and the universe. Integrating quantum computing as the engine for processing these enormous data streams, scientists can overcome the standard quantum limit, enabling detection of processes that are currently invisible to classical observation. The synergistic investment on quantum-native AI is essential for unraveling fundamental mysteries, from the Higgs boson's interactions to the origins of space-time.

Team quantum has an edge when playing games. Games have long been an important training ground for AI. In 1997, for instance, the Deep Blue supercomputer made history when it beat the human world champion in chess.[22] Throughout the 2010s, the Google DeepMind lab became known for training an AI to master Atari games by only giving it input to the pixels and the game score.[23] In 2016, DeepMind's AlphaGo computer program beat the world's highest ranked Go player.[24] Selecting the best move in a game is a form of optimization. Given the steady progress in quantum-enhanced optimization, it's plausible that in a future matchup, a Quantum AI would defeat a purely classical AI in games like chess or Go. Beyond optimization, quantum operations offer other intriguing advantages rooted in the fundamental physics of entanglement. Entanglement is a form of correlation. We say that two events are correlated when they co-occur together more often than one would expect from the frequency of their individual occurrence. This calculation employs standard probability theory, in which we use positive real-valued numbers in the interval of zero to one to describe the probability of an event occurring. The probabilities for different outcomes add up to one because by convention we assign a probability of one to an outcome that is certain to occur. Quantum physics can be seen as a generalization of standard probability theory, in which we assign probability amplitudes, complex-valued numbers, to the possible outcomes, and the sum of their probability amplitudes squared is normalized to one. Two states are entangled if they are correlated according to these rules of quantum probability. This is nicely explained in computer scientist Scott Aaronson's book *Quantum Computing Since Democritus*.[25] Two players who exchange entangled qubits can coordinate their moves in ways that are impossible with classical information alone, enabling coordination strategies that don't require direct communication or rely on classical information shared in real time. This has applications not only in games involving coordination between players, such as bridge or blackjack, but also in certain types of stock trades across geographically distant exchanges requiring synchronized actions without communication delays, as well as coordinated information processing in large-scale distributed neural networks.[26]

The trajectory of artificial intelligence, as highlighted by Turing Award laureate Richard Sutton in his influential 2019 essay "The Bitter Lesson," has largely been a story of escalating scale.[27] Progress has been consistently driven by employing ever-larger models and datasets, coupled with the application of general-purpose, distributed algorithms like gradient descent to train these models. This relentless pursuit of scale places enormous demands on computational resources, spurring significant investments in specialized hardware and even necessitating the fundamental redesign of data centers to manage these intensive workloads. Quantum processors offer an alternative computational paradigm that may help address this challenge. They operate by manipulating vectors within a mathematical realm known as Hilbert space. This space, central to quantum physics, represents all possible states of a quantum system. Hilbert space is vast; its vectors have an exponential number of elements. For a quantum computer with n qubits, the corresponding vector in Hilbert space can represent 2^n distinct values simultaneously. Remarkably, all these values can be processed or updated in a single operational step. However, a critical constraint exists: to control or read the vector, only a comparatively small, polynomial number of input and output bits can be exchanged with the quantum computer. This input/output bottleneck has led to the common perception that early quantum computers will be best suited for *small data, big compute* tasks.

Much less research has been devoted to figuring out how we might employ future large-scale quantum computers capable of handling the *big data* tasks underlying most of machine learning. But in a sign of things to come, a quantum algorithm has been proposed that can train certain wide and deep neural networks exponentially faster than classical methods.[28] This algorithm leverages quantum techniques for matrix operations, a core component of neural-network training. It achieves training times that scale logarithmically with the size of the training dataset, denoted as n. Such logarithmic scaling offers an exponential advantage over classical gradient descent. A crucial caveat, however, is the prerequisite of encoding the training data into a quantum state, a process requiring what is known as quantum random access memory (QRAM). Creating a QRAM state is an intensive step, taking order n operations. Yet once this quantum representation of the data is established, it can be used repeatedly to train multiple neural-network models with a significantly reduced computational cost, proportional to $log(n)$ (meaning the cost increases very slowly as the dataset grows) for each subsequent training. To illustrate the difference: training a neural network in a classical data center demands effort that scales as $t \cdot n^2$. This is because the number of training samples is typically chosen to be proportional to the number of weights in the neural network. Therefore, in each of the t training epochs we have to push n training examples through a neural network with n weights, hence the n^2 term. In contrast, the quantum approach would require an effort scaling as $n + t \cdot log(n)$. For datasets where n reaches into the trillions, the resultant savings in time and resources could be monumental. While the realization of large-scale QRAM is still aspirational, the cost and efficiency benefits it promises would ultimately favor quantum methodologies. One might envision a future when vast repositories of data and human knowledge are compactly encoded into quantum states, empowering quantum machine-learning algorithms to train models using a fraction of the current energy and hardware costs. We should note that it is still an open research question how QRAM that allows neural networks to perform a query in $log(n)$ time steps can be engineered in a fault-tolerant manner at a computational cost.[29]

There are additional quantum advantages relevant to the regime of large-scale computation and big-data tasks underpinning modern AI that do not require QRAM. Encoding classical data in quantum states can lead to exponential savings in communication for problems like inference with distributed graph neural networks and gradient-based training.[30] For some natural graph problems, quantum algorithms provide exponential memory savings over any classical algorithm.[31] Specifically, such discrete optimization problems and associated results in quantum communication complexity, which provide exponential savings in storage, suggest that more non-QRAM examples of quantum advantage exist for big data tasks. In many AI problems, memory, not computation, is the bottleneck. These types of advantages leverage the ability to represent information in quantum states that is in some sense a form of exponential compression yet crucially can still be used for computation. They make full use of the physical substrate of storage or communication that is allowed by quantum mechanics and that classical algorithms cannot take advantage of. A recent example is "quantum oracle sketching," which ingests massive sequences of classical data by mapping them onto small rotations of a quantum state. This state can then be used for classification or dimensionality reduction, offering exponential advantages over classical systems that lack the memory capacity to process such vast data streams.[32] Whenever computers communicate optically, they do not leverage the fact that one could instead send a quantum superposition of different configurations of light over the same communication channel, which can be used, for example, to solve classification problems that would otherwise require exponentially larger bandwidth. Large language models such as Google's Gemini require multiple distant data centers to communicate during training, and thus communication bottlenecks are a constraint on such computation that could be alleviated with the help of quantum computers.[33]

Recently, the term "agentic AI" has gained prominence, describing systems capable of effecting actions in the real world, such as booking a flight. This technical usage, however, is not intended to imply that artificial intelligence has agency or free will in a philosophical sense. Indeed, it is difficult to reconcile the concept of free will with a hardware substrate governed by the deterministic laws of classical mechanics and electrodynamics. Within such a framework, genuine choice and the emergence of agency – free will – become impossible to conceptualize. Everything proceeds on rails, and there is never truly a choice, a moment when agency could manifest. This does not fundamentally change when we add a source of randomness; if the decision between two alternatives is decided by the throw of a digital die, true agency remains absent.

Similarly, the computer science underlying today's AI has yet to provide a widely accepted explanation for what implements human consciousness. Humans have feelings and enjoy experiences; many of us would argue that this is what life is all about. Likewise, many of us have a sense of agency, believing we are at least partially in control of our destiny. Yet artificial intelligence is entirely stumped when challenged with explaining these experiential features associated with biological intelligence, a dilemma philosopher David Chalmers dubbed the "hard problem." To accurately predict the future of AI, we cannot dismiss these profound aspects of the human mind as mere epiphenomena and leave the discussion solely to philosophers.

Addressing these challenges, a recent article in *Entropy* coauthored by one of us puts forth a proposal drawing directly from the ontology of quantum mechanics, particularly aligning with the many-worlds formulation.[34] It suggests that consciousness is how we experience the emergence of a definite, classical world from the multitude of coexisting realities that, according to the many-worlds interpretation, the equations of quantum mechanics describe at any given moment. More specifically, the article suggests that a moment of consciousness arises whenever a quantum mechanical state enters into a superposition of two or more configurations. This moment of superposition, the authors argue in the article, is precisely the moment of choice when agency can manifest.

To our knowledge, this is the first proposal for the physical substance of consciousness that is amenable to experimental testing. This challenging but, in principle, feasible test consists of coherently linking quantum degrees of freedom in a human brain to a quantum computer, with the aim of investigating whether such a link allows for a controlled expansion of human conscious experience. To study how such a coherent link between qubits and neural tissue could be established, the article proposes experiments with xenon isotopes on brain organoids and fruit flies. Intriguingly, early (still unpublished) experimental data show that different xenon isotopes exhibit different anesthetic potencies. If this isotope effect is confirmed, then it would provide strong evidence that quantum processes need to be considered when modeling the action of anesthetics or more generally when formulating a theory of consciousness.

Quantum generalizations have been proposed for most classical machine-learning methods. Take any acronym for a popular machine-learning (ML) method, put a Q in front of it, and a Google search will point you to a paper describing the quantum version (QML). Often such generalizations come with a theoretical argument about why they should be superior. For instance, quantum kernel methods can access large feature spaces inaccessible to their classical counterparts, quantum GANs (Generative Adversarial Networks) or quantum diffusion networks can work with probability distributions that can't be represented on classical computers. But do such theoretical advantages really matter in machine-learning practice? We don't know yet. This is because machine learning is manifestly a heuristic discipline. Whether a task is learnable or whether the architecture of a neural network is well suited for a learning task is typically found by trial and error. Indeed, classical neural networks were discovered experimentally, and it remains a major open research problem to analytically and rigorously explain why they are so effective. Similarly, it remains to be empirically tested whether the quantum generalizations of machine learning deliver better results. But the testing of heuristic QML algorithms will have to wait until quantum computers reach the size at which they can handle datasets big enough to be of practical relevance. Table 1 explores why different QML algorithms may have an advantage.

While most QML algorithms today have only theoretical advantages, there exists an early empirical foundation built on smaller-scale quantum processors: early Nuclear Magnetic Resonance platforms provided experimental proof-of-principle demonstrations for classifying quantum states and testing elementary QML protocols, bridging the gap between theoretical algorithms and physical realization. As hardware scales, these early results provide confidence that the theoretical speedups will manifest empirically.

While quantum computing holds promise for advancing artificial intelligence in the future, AI is already accelerating quantum technology today. This synergy is likely to create a virtuous cycle, leading to the emergence of new capabilities. Where is AI boosting quantum computing and where might its impact be greatest?

Quantum error correction – the technology that ensures that quantum information does not leak into the environment and that all operations necessary to execute a quantum algorithm are executed with high precision – is crucial for scalable quantum computing, but the hardware overhead to achieve it remains high. Machine learning is being deployed to discover more efficient methods for encoding quantum information stably and for decoding the error syndromes reported by

Table 1
Advantages of Various Quantum Machine Learning Algorithms

Classical AI/ ML Task	Underlying Computation	Quantum Algorithm(s)	Potential Quantum Advantage	Challenges/ Caveats
Model training	Parameter optimization	QAOA, VQE, DQI, quantum Gibbs sampling, quantum streaming	Faster convergence, escape local minima, find local minima	Superquadratic speedups are only established for very specific optimization problems, less memory
Generative modeling	Sampling complex distributions	QBMs, QGANs, quantum noise in diffusion networks	Higher expressivity, better sample quality	Practical advantage in generating compelling images, videos, or other patterns not yet established
Classification	Kernel computation	QSVM (quantum kernels)	Access intractable feature spaces	Demonstrating advantage on real data, data encoding
Linear regression	Matrix inversion $Ax = b$	HHL algorithm	Exponential speedup in N (dimension)	Data I/O overhead, sensitivity to properties of the Matrix A
Dimensionality reduction	Eigenvalue decomposition	QPCA (quantum phase estimation)	Exponential speedup in N (dimension)	Data I/O overhead
Data search	Unstructured search	Grover's algorithm	Quadratic speedup	Oracle implementation efficiency, quadratic speedup requires long runtimes to manifest

Classical AI/ ML Task	Underlying Computation	Quantum Algorithm(s)	Potential Quantum Advantage	Challenges/ Caveats
Pattern matching	Finding hidden shifts	Montanaro, Kuperberg	Superpolynomial speedup	Unclear whether this is a naturally occurring problem
Anomaly detection	Search for rare events	Grover's algorithm	Quadratic speedup	Oracle implementation efficiency, quadratic speedup requires long runtimes to manifest

Table 1 shows how quantum machine-learning algorithms have been suggested for most classical machine-learning tasks. The quantum algorithms abbreviated above are QAOA (Quantum Approximate Optimization Algorithm), VQE (Variational Quantum Eigensolver), DQI (Decoded Quantum Interferometry), QBMs (Quantum Boltzmann Machines), QGANs (Quantum Generative Adversarial Networks), QSVM (Quantum Support Vector Machine), HHL (Harrow–Hassidim–Lloyd), and QPCA (Quantum Principal Component Analysis). Source: Authors' data and analysis.

qubit measurements designed to detect errors. A notable early success in this area is AlphaQubit, developed through a collaboration between the Google DeepMind and Quantum AI teams. This ML-based decoding system maps error syndromes to suspected errors. This task is well suited for transformer networks, the neural network architecture underlying large language models, because we can generate abundant synthetic training data by injecting errors into simulated circuits and measuring the resulting syndromes, which then enables us to fine-tune the network with experimental training data from actual hardware runs. The result was that AlphaQubit outperformed the best human-designed methods in error rate.[35]

AI's recent advancements in generating and refining programs for complex problems – exemplified by DeepMind's AlphaEvolve tackling classical optimization challenges – suggest a powerful synergy with quantum computing.[36] By expanding the search space to include quantum programs, the AlphaEvolve coding agent could be employed to discover efficient quantum circuits for classically intractable problems.

Beyond software and algorithms, AI's potential extends to the processes of designing and refining the quantum hardware itself, further fueling the virtuous cycle. For instance, in chip design, machine-learning models can be trained on extensive databases of existing designs and their corresponding performance metrics. The Google Quantum AI team is actively exploring this, leveraging metrology data to teach AI how design parameters map to chip performance, with the goal of eventually enabling automated and optimized chip design based on a sufficiently large and comprehensive dataset.

Beyond the initial design phase, processor calibration presents another challenge for which artificial intelligence excels. Tuning a novel quantum processor to its optimal operational regime involves navigating a high-dimensional parameter space – a task akin to finely tuning a complex musical instrument. Here, sophisticated AI algorithms are indispensable for efficiently solving these intricate optimization problems and achieving peak performance from the quantum hardware. Furthermore, maintaining the quality and integrity of quantum hardware throughout the manufacturing process significantly benefits from AI-driven wafer inspection. Computer vision techniques, powered by machine learning, automate and enhance the classification of wafer and circuit integrity. This application is yet another example of how AI accelerates a critical step in the hardware development pipeline.

The fundamental ways by which quantum computing and artificial intelligence can help each other are hardware-agnostic and apply to all scalable approaches to quantum computing, including those based on superconducting qubits, neutral atoms, trapped ions, photons, spin qubits, or topological qubits.

To predict how quantum computing and AI will mutually enhance and evolve toward Quantum AI, it's instructive to recall AI's history (Figure 2): Following initial conceptual breakthroughs, significant progress has largely been driven by the scaling of underlying hardware. This scaling enables the creation of larger models and the ingestion of the vast datasets essential for their training. The perceptron, the first multilayer neural network and often considered the precursor to modern neural networks, was proposed by psychologist Frank Rosenblatt in 1957. The architecture of a modern deep neural network powering a large language model is not fundamentally different. Both consist of layers of McCulloch-Pitts neurons – a simplified computational model of a biological neuron that receives inputs, multiplies each by a corresponding weight, sums these weighted inputs, and generates an output if this sum exceeds a predefined threshold. A deep neural network is trained with the backpropagation algorithm, a development that started in the 1970s. However, it wasn't until the mid-1980s, particularly through the joint work of psychologist David Rumelhart and computer scientists Geoffrey Hinton and Ronald J. Williams, that its significance for training artificial neural networks was fully realized, becoming a cornerstone of the field. The reason we are seeing an explosion in the performance and use of neural networks is that the underlying classical hardware has become exponentially more performant. Moore's law once

Figure 2
Historic Milestones in Quantum Computing and AI

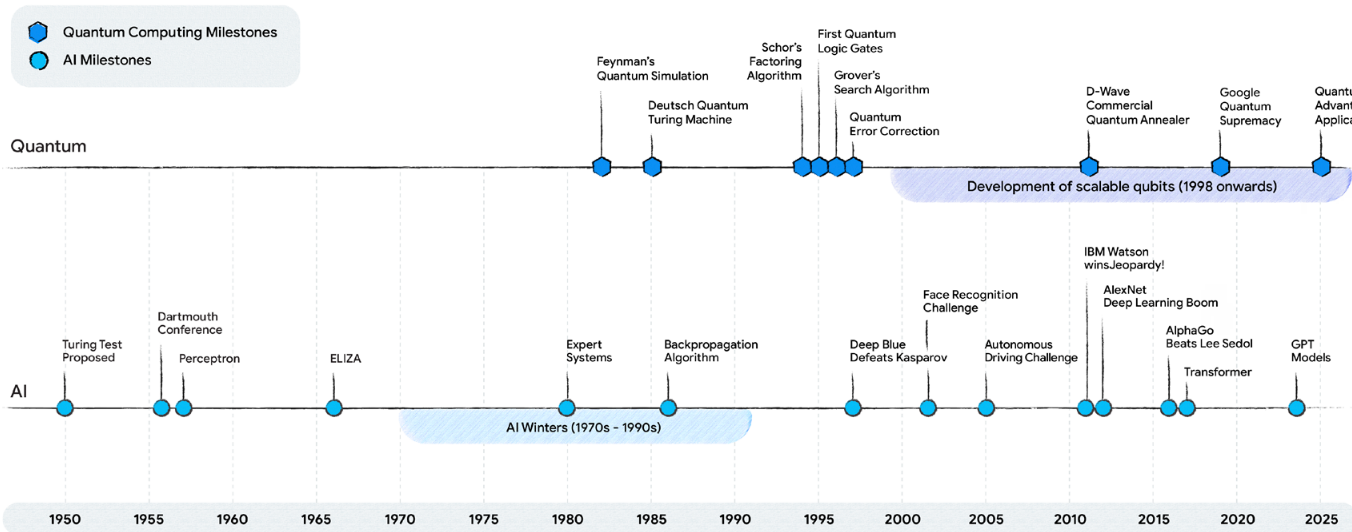

Figure 2 shows seminal events in the histories of quantum computing and AI. AI is about seventy-five years in the making; quantum computing is only about half as old. Our prediction is that well before AI turns one hundred, it will merge with quantum computing and morph into Quantum AI. Source: Figure by the authors.

ensured that the power of an individual processor doubled about every two years. As that progress began to slow, classical parallel processing compensated, leading to the current era in which massive data centers are filled with processors, such as graphics processing units (GPUs) and tensor processing units (TPUs), specialized for the tensor computations needed in neural network training and operation. This immense compute power, coupled with vast multimodal datasets, unlocked the potential of long-established AI concepts. Computational innovations such as the attention mechanism of transformer networks that allowed parallelization of the training for language models were certainly important, but the fundamental driver enabling the AI revolution has been hardware progress and large datasets. The machine-learning algorithms have remained remarkably unchanged.

Now the hardware powering artificial intelligence is once again running up against limitations. Besides the costs of so many TPUs or GPUs, the energy need is becoming substantial. Generating a single image uses about 10 watt-hours, roughly the amount of energy to charge a cell phone; training a large language model requires about 10 gigawatt-hours, an amount a nuclear power plant outputs in ten hours. Can we expect that quantum hardware will help to alleviate these hurdles and to ensure the continued progress of AI? Because of the new capabilities afforded by quantum computers discussed in this essay, the answer seems to be "yes." The fact that vast amounts of training data can be stored and manipulated succinctly in superposition is a key reason why quantum resources may one day curb the steeply rising costs of training large neural networks.

But beyond these practical efficiencies, there's a more fundamental argument: Since information must be represented by configurations of matter, it is ultimately the laws of physics that dictate which operations are possible within any information-processing system. It follows that the most advanced AI systems will eventually leverage the richest computational framework allowed by physics, perhaps eventually even incorporating principles from yet undiscovered physics in quantum gravity and elsewhere, should such theories yield more powerful computational models. From this perspective, continued progress in AI seems to lead inexorably toward Quantum AI. Hence, we dare to predict that by the time artificial intelligence turns one hundred years old, it will have morphed into Quantum AI.

ABOUT THE AUTHORS

Maria Spiropulu is a particle physicist and the Shang-Yi Ch'en Professor of Physics at the California Institute of Technology. She is Cofounder of Bohr Quantum Technology, an advanced quantum networking technologies startup. Her recent research includes testing wormhole teleportation protocols, advancing quantum teleportation networks, optimizing real time AI on field-programmable gate arrays and application-specific integrated circuits, and developing and building an ultraprecision multichannel timing sensor system for the High-Luminosity Large Hadron Collider at CERN. She was Cochair of *Elementary Particle Physics: The Higgs and Beyond* 2025, a consensus study report for the National Academies of Sciences, Engineering, and Medicine that presents a forty-year strategic vision for the field of particle physics. She also coauthored the Department of Energy's 2020 *Quantum Internet Blueprint* report.

Hartmut Neven is Vice President of Engineering at Google, where he leads the Quantum Artificial Intelligence Lab, which he founded in 2012. The Google Quantum AI team has achieved a number of firsts, including implementing the first quantum computations beyond the reach of classical supercomputers, demonstrating quantum error correction, and executing the first verifiable quantum computation.

ENDNOTES

[1] Wikipedia Contributors, "APX," Wikipedia, last edited March 24, 2025, https://en.wikipedia.org/wiki/APX.

[2] Stephen P. Jordan, Noah Shutty, Mary Wootters, et al., "Optimization by Decoded Quantum Interferometry," arXiv (2024), https://arxiv.org/abs/2408.08292; and Edward Farhi, Sam Gutmann, Daniel Ranard, and Benjamin Villalonga, "Lower Bounding the MaxCut of High Girth 3-Regular Graphs Using the QAOA," arXiv (2025), https://doi.org/10.48550/arXiv.2503.12789.

[3] Stephen P. Jordan and Google Quantum AI, "Quantum Algorithm Zoo," last updated October 30, 2025, https://quantumalgorithmzoo.org.

[4] "XPRIZE Quantum Applications," XPRIZE Foundation, https://www.xprize.org/prizes/qc-apps (accessed May 2, 2025).

[5] Stefan Chmiela, Valentin Vassilev-Galindo, Oliver T. Unke, et al., "Accurate Global Machine Learning Force Fields for Molecules with Hundreds of Atoms," *Science Advances* 9 (2) (2023), https://doi.org/10.1126/sciadv.adf0873.

[6] Ibid.; and James Kirkpatrick, Brendan McMorrow, David H. P. Turban, et al., "Pushing the Frontiers of Density Functionals by Solving the Fractional Electron Problem," *Science* 374 (6573) (2021): 1385–1389.

[7] John Jumper, Richard Evans, Alexander Pritzel, et al., "Highly Accurate Protein Structure Prediction with AlphaFold," *Nature* 596 (7873) (2021): 583–589; and Kirkpatrick, McMorrow, Turban, et al., "Pushing the Frontiers of Density Functionals by Solving the Fractional Electron Problem."

[8] Dmitry A. Abanin, Rajeev Acharya, Laleh Aghababaie-Beni, et al., "Constructive Interference at the Edge of Quantum Ergodic Dynamics," arXiv (2025), https://doi.org/10.48550/arXiv.2506.10191; and C. Zhang, R. G. Cortiñas, A. H. Karamlou, et al., "Quantum

Computation of Molecular Geometry via Many-Body Nuclear Spin Echoes," arXiv (2025), https://doi.org/10.48550/arXiv.2510.19550.

9 Nicholas C. Rubin, Dominic W. Berry, Alina Kononov, et al., "Quantum Computation of Stopping Power for Inertial Fusion Target Design," *Proceedings of the National Academy of Sciences* 121 (23) (2024), https://doi.org/10.1073/pnas.2317772121.

10 Robin Kothari and Rolando Somma, "A New Quantum Algorithm for Classical Mechanics with an Exponential Speedup," The Google Research Blog, December 4, 2023, https://research.google/blog/a-new-quantum-algorithm-for-classical-mechanics-with-an-exponential-speedup; and Aram W. Harrow, Avinatan Hassidim, and Seth Lloyd, "Quantum Algorithm for Linear Systems of Equations," *Physical Review Letters* 103 (15) (2009): 150502.

11 Juan Maldacena and Leonard Susskind, "Cool Horizons for Entangled Black Holes," *Fortschritte der Physik* 61 (9) (2013): 781–811, https://doi.org/10.1002/prop.201300020; and Ping Gao, Daniel Louis Jafferis, and Aron C. Wall, "Traversable Wormholes via a Double Trace Deformation," *Journal of High Energy Physics* 2017 (12) (2017): 1–25.

12 Daniel Jafferis, Alexander Zlokapa, Joseph D. Lykken, et al., "Traversable Wormhole Dynamics on a Quantum Processor," *Nature* 612 (7938) (2022): 51–55.

13 Juan Maldacena and Douglas Stanford, "Remarks on the Sachdev-Ye-Kitaev Model," *Physical Review D* 94 (10) (2016): 106002.

14 Garnet Kin-Lic Chan and Sandeep Sharma, "The Density Matrix Renormalization Group in Quantum Chemistry," *Annual Review of Physical Chemistry* 62 (2011): 465–481.

15 Srihari S. Iyengar, "Quantum Chemical Dynamics at the Nanoscale: The Role of Nuclear Quantum Effects, Non-Adiabaticity, and Time-Dependence," *Chemical Reviews* 123 (1) (2023): 405–466.

16 Hsin-Yuan Huang, Michael Broughton, Jordan Cotler, et al., "Quantum Advantage in Learning from Experiments," *Science* 376 (6598) (2022): 1182–1186.

17 Google Research, "Quantum Transduction and Networking for Scalable Computing Applications," https://research.google/programs-and-events/quantum-transduction-and-networking-for-scalable-computing-applications (accessed May 3, 2025).

18 Richard R. Allen, Francisco Machado, Isaac L. Chuang, et al., "Quantum Computing Enhanced Sensing," arXiv (2025), https://doi.org/10.48550/arXiv.2501.07625.

19 Maldacena and Stanford, "Remarks on the Sachdev-Ye-Kitaev Model."

20 Akash V. Dixit, Srivatsan Chakram, Kevin He, et al., "Searching for Dark Matter with a Superconducting Qubit," *Physical Review Letters* 126 (14) (2021): 141302.

21 Aaron Chou, Kent Irwin, Reina H. Maruyama, et al., "Quantum Sensors for High Energy Physics," arXiv (2023), https://doi.org/10.48550/arXiv.2311.01930.

22 Feng-Hsiung Hsu, *Behind Deep Blue: Building the Computer That Defeated the World Chess Champion* (Princeton University Press, 2022).

23 Volodymyr Mnih, Koray Kavukcuoglu, David Silver, et al., "Human-Level Control Through Deep Reinforcement Learning," *Nature* 518 (7540) (2015): 529–533.

24 David Silver, Aja Huang, Chris J. Maddison, et al., "Mastering the Game of Go with Deep Neural Networks and Tree Search," *Nature* 529 (7587) (2016): 484–489.

25 Scott Aaronson, *Quantum Computing Since Democritus* (Cambridge University Press, 2013).

[26] Sadiq Muhammad, Armin Tavakoli, Maciej Kurant, et al., "Quantum Bidding in Bridge," *Physical Review X* 4 (2) (2014): 021047; Joseph X. Lin, Joseph A. Formaggio, Aram W. Harrow, and Anand V. Natarajan, "Quantum Blackjack or Can MIT Bring Down the House Again?" arXiv (2019), https://arxiv.org/abs/1908.09417; Dawei Ding and Liang Jiang, "Coordinating Decisions via Quantum Telepathy," arXiv (2024), https://doi.org/10.48550/arXiv.2407.21723; and Dar Gilboa, Hagay Michaeli, Daniel Soudry, and Jarrod R. McClean, "Exponential Quantum Communication Advantage in Distributed Inference and Learning," in *NIPS '24: Proceedings of the 38th International Conference on Neural Information Processing Systems*, ed. Amir Globerson, Lester Mackey, Danielle Belgrave, et al. (Curran Associates, Inc., 2024), 30425–30473.

[27] Richard Sutton, "The Bitter Lesson," March 13, 2019, http://www.incompleteideas.net/IncIdeas/BitterLesson.html; and Jared Kaplan, Sam McCandlish, Tom Henighan, et al., "Scaling Laws for Neural Language Models," arXiv (2020), https://doi.org/10.48550/arXiv.2001.08361.

[28] Alexander Zlokapa, Hartmut Neven, and Seth Lloyd, "A Quantum Algorithm for Training Wide and Deep Classical Neural Networks," arXiv (2021), https://doi.org/10.48550/arXiv.2107.09200.

[29] Vittorio Giovannetti, Seth Lloyd, and Lorenzo Maccone, "Quantum Random Access Memory," *Physical Review Letters* 100 (16) (2008): 160501; and Alexander M. Dalzell, András Gilyén, Connor T. Hann, et al., "A Distillation-Teleportation Protocol for Fault-Tolerant QRAM," arXiv (2025), https://doi.org/10.48550/arXiv.2505.20265.

[30] Gilboa, Michaeli, Soudry, and McClean, "Exponential Quantum Communication Advantage in Distributed Inference and Learning."

[31] Chi-Fang Chen, Hsin-Yuan Huang, John Preskill, and Leo Zhou, "Local Minima in Quantum Systems," in *STOC 2024: Proceedings of the 56th Annual ACM Symposium on Theory of Computing* (Association for Computing Machinery, 2024), 1323–1330.

[32] Haimeng Zhao, Alexander Zlokapa, Hartmut Neven, et al., "Exponential Quantum Advantage in Processing Massive Classical Data," arXiv (2026), https://arxiv.org/pdf/2604.07639.

[33] John Kallaugher, Ojas Parekh, and Nadezhda Voronova, "Exponential Quantum Space Advantage for Approximating Maximum Directed Cut in the Streaming Model," in *STOC 2024: Proceedings of the 56th Annual ACM Symposium on Theory of Computing* (Association for Computing Machinery, 2024), 1805–1815; Seid Koudia, Angela Sara Cacciapuoti, Kyrylo Simonov, and Marcello Caleffi, "How Deep the Theory of Quantum Communications Goes: Superadditivity, Superactivation and Causal Activation," *IEEE Communications Surveys & Tutorials* 24 (4) (2022): 1926–1956; and Gemini Team Google, "Gemini: A Family of Highly Capable Multimodal Models," arXiv (2023), https://doi.org/10.48550/arXiv.2312.11805.

[34] Hartmut Neven, Adam Zalcman, Peter Read, et al., "Testing the Conjecture That Quantum Processes Create Conscious Experience," *Entropy* 26 (6) (2024): 460.

[35] Johannes Bausch, Andrew W. Senior, Francisco J. H. Heras, et al., "Learning High-Accuracy Error Decoding for Quantum Processors," *Nature* 635 (8040) (2024): 834–840, https://doi.org/10.1038/s41586-024-08148-8.

[36] Bernardino Romera-Paredes, Mohammadamin Barekatain, Alexander Novikov, et al., "Mathematical Discoveries from Program Search with Large Language Models," *Nature* 625 (7995) (2024): 468–475.