Artificial intelligence and quantum computing white paper

March 1st, 2025

Table of contents

1 - Executive summary

Two computing revolutions are currently in the making with artificial intelligence (AI) and quantum computing, with different levels of maturity and market footprints. While the European Union's scientific position in both domains is significant, it is still squeezed between the USA dominance and China's increasing role, particularly with artificial intelligence. Just in the large language models, the emergence of DeepSeek has shown that innovation, scientific astuteness and open source models can significantly alter the market balance.

Leadership in these domains comes from scientific excellence and the capability to create strong integrated software platforms, build significant, scalable and efficient computing infrastructure and spur the creation of as generic as possible use cases for end users from the enterprize and public services to the consumer markets. A tight integration between academic research and industry R&D is a fundamental enabler of market success.

Quantum computing is still a promise in the making but its synergies with AI are already there and growing. Many AI-driven techniques are already enabling significant progress in quantum computing research and industry developments, from optimizing qubit control and quantum error mitigation (QEM) strategies in the early stage noisy intermediate-scale quantum (NISQ) regime to designing novel quantum algorithms, highlighting the importance of mastering this dual discipline as an enabling technology.

Likewise and on the other way around, early evidence of quantum advantage in faster computing times, better results or the need for less training data for solving specific machine learning problems illustrates how uniting AI and quantum resources could bring value, even though the required large-scale, fault-tolerant quantum computers (FTQC) needed to obtain such advantage remain a longer-term objective.

Public investment in the convergence of AI and quantum computing could strengthen the EU competitiveness by building on the existing research excellence in both fields and accelerate the transition from laboratories to market applications. One strategic goal of employing quantum computing is to advance AI-based solutions for healthcare, finance, materials discovery, and security. It is a mid to long-term target, but one that will determine industrial and scientific leadership in both domains. Likewise, the EU AI and quantum research landscape should encourage the development of both open source and commercial integrated software engineering platforms.

As international competition escalates, ensuring support for academic research and private-sector innovation at this intersection will not only secure the economic benefits associated with disruptive breakthroughs but also reinforce the EU's position as a prominent global player in emerging deep-tech ecosystems. A carefully orchestrated funding strategy that spans fundamental research, talent cultivation, and technology transfer incentives will ensure a robust pathway from visionary lab-scale projects to tangible, high-impact industry platforms and applications. Accordingly, this white paper lays out the scientific and applications landscape for consolidating artificial intelligence and quantum computing disciplines with providing a research and use case agenda.

It starts with describing how quantum computing could help develop innovative artificial intelligence solutions, particularly in the machine learning spaces. This is a mid to long term ploy. It is aligned with quantum computer hardware roadmaps.

It then covers the use cases of classical AI to empower research and developments of quantum technologies, focused on quantum computing and quantum sensing. This application domain of AI will mature. One important aspect is to ensure classical artificial intelligence scales well as the requirements of quantum computing platforms will grow, as the domain progressively shifts from NISQ to FTQC quantum computers. One example lies with the critical role of machine learning empower quantum error correction (QEC) techniques. At last, it provides a longer term research agenda to drive work in foundational questions related to how AI and quantum computing interact and benefit each other.

The white paper ends with a set of recommendations and challenges on the way to orchestrate the proposed theoretical work, align quantum artificial intelligence developments with quantum hardware roadmaps, work on both classical and quantum resource estimates, particularly with the goal to mitigate and optimize energy consumption, orchestrate this upcoming hybrid software engineering discipline and develop the European industry competitiveness while considering societal aspects.

2 - Introduction

The convergence of artificial intelligence (AI) and quantum computing is a rapidly evolving field with the potential to impact numerous aspects of technology and science. This white paper explores the synergistic relationship between these two disciplines, outlining how quantum computing can enhance classical AI capabilities and how AI can be used to advance quantum computing. It details different methods, applications, and objectives, with a focus on both immediate and long-term goals. At this stage, it can be viewed as an early proposal for a strategic and industry research roadmap.

This document has been written by AI and quantum technologies (QT) experts in the spirit of the European Strategic Research and Innovation Agenda (SRIA), that is trying to avoid overselling. Nevertheless, some parts of the text are speculative, due to the great novelty of the subject.

One of the key areas of focus is quantum-assisted machine learning, where quantum processors pre-process classical data that then feed classical AI. This approach may lead to improvements in total processing speed, accuracy, and reduce the amount of training data required. Quantum data pre-processing, where, for example, quantum simulators feed data into classical AI algorithms, is such a technique. Quantum computing is also explored to accelerate the training phase of classical machine learning models, using both near-term noisy intermediate-scale quantum (NISQ) devices with variational algorithms, and future fault-tolerant quantum computers (FTQC). Another approach is learning with quantum models, where quantum computing takes the helm for both the training and inference phases, potentially uncovering data patterns that are intractable for classical systems.

Quantum computing could significantly improve reinforcement learning (RL), addressing computational bottlenecks and lengthy training times. Quantum reinforcement learning (QRL) methods can use parameterized quantum circuits to optimize decision-making in complex environments, particularly in industrial applications. In unsupervised learning, quantum algorithms are developed to handle tasks like automatic clustering and dimensionality reduction, with the potential to provide exponential or polynomial speedups compared to the best classical methods available. This can for example be useful to develop innovative cyberthreat detection solutions.

These combined technologies have broad potential applications. In healthcare and life sciences, quantum simulations may be used to generate training data for AI models, which could accelerate drug discovery by exploring new chemical spaces. Medical image analysis used with X-rays, MRI scans and biological sample imaging could be enhanced, reducing reliance on large and labeled datasets. In time series analysis, quantum methods may provide more efficient modeling and faster anomaly detection. Quantum computing is also being explored for extracting insights from complex quantum systems.

AI is also being used to advance quantum computing through the design of novel quantum algorithms and protocols, the optimization of quantum circuits, with quantum error mitigation, and the implementation of quantum error correction, particularly for the costly task of error syndrome detection. AI could help in discovering and optimizing quantum experiments, simulating quantum systems, and analyzing quantum data. We can even envision to creating fully quantum AI models, where all data, training algorithms, and inference systems are quantum in nature.

The field faces several challenges, including current hardware limitations such as qubit numbers, fidelities, and scalability. There are difficulties in loading and processing classical data into quantum states at various levels. Training quantum models also presents challenges such as barren plateaus and the lack of efficient quantum equivalents to classical back propagation used in the training of neural networks. Additionally, ensuring trust, robustness, interpretability, and explainability of AI models as well as avoiding various data biases, are critical for the reliable application of these technologies in practical situations. Standardized interfaces should be developed to share data and translate quantum problems into a common machine learning language.

This white paper outlines both short-term research goals (3 to 5 years), mid-term research goals (5 to 10 years) and long-term research goals (beyond 10 years) goals related to these various challenges. For example, in the short term, the focus is on demonstrating quantum utility for chemistry and error mitigation, identifying features that are easier to extract using quantum machines, and using AI to rediscover known quantum algorithms.

Medium-term objectives include establishing hybrid frameworks for molecular simulations, new materials development, scaling hybrid classical-quantum models and improving quantum error correction techniques. Long-term goals involve validating frameworks for drug retargeting tasks, broadening the scope of quantum-enhanced machine learning, developing fully quantum AI models, and using AI to co-design quantum algorithms and quantum hardware.

Finally, the white paper also emphasizes the importance of foundational questions about learning in a quantum world. These include understanding the connections between physics and quantum machine learning, defining learning with quantum data, and exploring the limitations of quantum autonomous agents. The need for interdisciplinary collaboration, open-source software, and standardized data sets is also highlighted, along with the need to train the next generation of experts in both quantum information science and machine learning.

In conclusion, the combination of quantum computing and AI has the potential to drive significant advances across many sectors. Strategic research, addressing key challenges, and fostering collaboration will be of paramount importance for realizing the full potential of these technologies to the benefit of society and European competitiveness in both the AI and quantum computing fields.

3 - Quantum for AI

The rapid advancements in artificial intelligence across scientific and industrial domains have underscored the need to overcome the computational limitations of classical methods and explore alternative paradigms for scalable and efficient AI solutions. Recent progress in quantum computing suggests the potential for quantum-enhanced AI approaches to outperform purely classical techniques, particularly in addressing computationally intensive tasks. A key direction in this integration is the development of hybrid quantum-classical architectures, where quantum processors serve as pre-processing units for classical AI inference tasks. In the near to mid-term, this integration is expected to be feasible with noisy intermediate-scale quantum devices comprising 100 to 200 physical qubits, while in the long term, early-stage fault-tolerant quantum computers with over 50 logical qubits could enable more complex algorithms.

Quantum-enhanced AI primarily revolves around accelerating specific subroutines, such as optimization, sampling, and high-dimensional data processing, which are computationally expensive for classical methods. One proposed approach involves combining quantum processors with high-performance computing (HPC) resources into hybrid systems that leverage quantum algorithms for specific computational bottlenecks while retaining classical AI's robustness and scalability. Another promising avenue is the use of quantum-generated data to enhance AI models, potentially improving processing speed, computational complexity, modeling accuracy, and the amount of data required for training. While these developments highlight the transformative potential of quantum-assisted AI, realizing practical impact requires sustained interdisciplinary efforts from both the classical AI and quantum computing communities. This section briefly delineates the current state of quantum computing methods and their integration with classical workflows, with an emphasis on machine learning techniques. It further identifies key AI subfields—such as optimization, multi-agent systems, and reasoning—that remain underexplored despite the significant promise of quantum technologies, thereby underscoring the need for further research to facilitate broad and effective adoption.

The goals are the following:

- Demonstrate quantum utility from using quantum processors as a pre-processing stage for classical AI inference tasks or full end-to-end quantum machine learning solutions.
- Demonstrate this at scales achievable in the near- to mid-term with 100 to 200 physical qubits using quantum error mitigation and variational circuits (NISQ) or in the long-term early stage fault-tolerant quantum computers supporting between 50 and 100 logical qubits which require a much larger number of physical qubits and allow much deeper quantum circuits.

The proposed scenarios are:

- Quantum processors and HPC combined into hybrid systems, including hybrid algorithms.
- AI using data produced from a quantum processor, achieving overall improvement in
    - processing speed and computational complexity,
    - modeling capability and response accuracy,
    - number of samples needed for training.
- Classical and quantum resource estimation to identify expected usefulness and timeline.

3.1 - Quantum-assisted machine learning

3.1.1 - Supervised learning

Quantum supervised learning (QSL) refers to the application of quantum algorithms to solve supervised learning tasks. Supervised learning is a fundamental area of machine learning where a model is trained on labeled data to learn the relationship between inputs (features) and corresponding outputs (labels). The goal is to generalize this learned relationship to predict the outputs for unseen inputs. Common examples of supervised learning tasks include classification, where the model assigns an input to a specific category, and regression, where the model predicts a continuous quantitative value.

State-of-the-art methods, such as neural networks, have demonstrated remarkable performance in many tasks, particularly when large datasets and specialized hardware are available. However, these methods face several challenges, including the need for extensive sets of labeled data, long training times for complex models, limited explainability of predictions, and the high computational cost of modeling high-dimensional, non-linear relationships. These limitations often hinder the scalability and practicality of classical approaches in real-world scenarios.

Examples include expert annotation costs in fields like medical imaging and legal analysis that limit scalability, while noisy labels from crowd-sourcing or subjective tasks, like sentiment analysis, hinder performance and generalization. QSL aims to address these challenges, tackling problems that are difficult or impossible for classical methods. In general, using quantum algorithms for supervised learning problems has different scopes: to simplify the training procedure in terms of training time and the number of optimization cycles required; to improve performance in terms of accuracy - how well the models learn general patterns from the data; to enhance efficiency by reducing the amount of data needed for training, striving to achieve comparable or superior results with less information. This field encompasses a range of approaches that consider integrating quantum technologies into machine learning workflows, categorized by the extent of quantum involvement and the type of quantum hardware used. The distinction between FTQC and NISQ technologies, where fault tolerance is not assumed to be available, shapes the methodologies and practical applications of QSL, influencing both their theoretical potential and real-world feasibility.

Efficient quantum training of classical models

One key methodological approach of quantum computing in supervised learning is its use to enhance the training phase of classical models. Fault-tolerant quantum machine learning aims to achieve a theoretical speedup in optimizing well-defined classical algorithms, such as support vector machines, splines, and linear regression by leveraging quantum algorithms to solve the underlying parametric optimization problems more efficiently.

These methods assume a specific functional form for the target function of interest, often expressed as a linear relationship with parameters to be estimated, such as in least squares regression, thus enabling the use of quantum techniques like variants of the Harrow-Hassidim-Lloyd (HHL) algorithm to accelerate computationally expensive linear algebra operations. These methods depend on the availability of error-corrected qubits and long quantum circuits, which enable efficient solutions to convex optimization problems with polynomial complexity. Strong guarantees for the performance of quantum training of classical models have in particular been found for models which are both sufficiently dissipative and sparse, as they arise in the context of quantum training pruned classical networks.

In parallel, NISQ-based approaches, such as shadow models, aim to leverage the current generation of quantum hardware to train classical surrogates of quantum neural networks that mimic quantum models under the assumption that training of these can be performed more efficiently.

Additionally, in the neural network based deep learning realm, classical-inspired quantum techniques, like quantum convolutional neural networks (QCNNs) or quantum perceptrons, adapt established classical architectures to leverage computational advantages during training. All these approaches confine the use of quantum resources to the training phase, enabling classical inference for scalability and broader end-user applicability. However, in these cases, ad hoc techniques must be employed to ensure that the quantum component is used only during training, while keeping the algorithm entirely classical during testing. One such approach is knowledge distillation, where a quantum-trained model guides the training of a purely classical model. To transfer knowledge from a quantum neural network (QNN) to a classical neural network, first, a variational quantum circuit based QNN is trained on a classification or regression task, obtaining softmax-like output probabilities. Next, a classical neural network is trained using the soft labels predicted by the QNN instead of the ground truth. The goal is to generate a classical model that mimics the QNN’s decision boundary, effectively capturing its learned representations. This approach is particularly useful when classical models struggle with optimization, generalization, or data efficiency.

For instance, in image classification tasks, deep classical networks often require large labeled datasets and extensive hyperparameter tuning to achieve efficient convergence. Similarly, in time series forecasting, classical models face challenges with long-range dependencies and high-dimensional correlations, necessitating extensive feature engineering and computationally expensive training. In molecular property prediction, classical models may struggle to capture quantum mechanical properties using traditional feature representations, leading to suboptimal performance.

Learning with quantum models

Beyond training enhancements, learning with quantum models incorporates quantum computing into both the training and inference phases, aiming to uncover patterns in data that are intractable for classical systems. This involves leveraging complex quantum kernels and parameterized quantum circuits to represent data in ways classical models cannot. The variational paradigm, which relies on hybrid quantum-classical techniques in which a smaller variational quantum circuit is controlled by a classical algorithm, is particularly prominent in NISQ-based implementations, offering opportunities to explore novel hypothesis functions.

Preliminary results indicate the potential of quantum models to reduce the parameter space, require less training data, and enable more efficient training procedures. These results currently hold only for specific scenarios involving quantum data. Investigating these advantages in data-intensive domains, such as genomics (e.g., DNA sequence analysis), financial modeling (e.g., high-frequency trading risk assessment), climate modeling (e.g., large-scale weather prediction), and industrial Internet of Things (e.g., real-time sensor network optimization), could open the possibility for quantum models to replace classical ones in both training and inference.

However, significant challenges persist, including barren plateaus in optimization and the lack of efficient quantum equivalents to the classical back propagation training technique. In some instances, the quantum models can even be fully dequantized, referring to the situation in which an efficient classical algorithm for the same task can be found, or classical surrogates formulated. At the same time, one has to ensure that the quantum circuit is sufficiently expressive. Despite these obstacles, the integration of quantum computing across the supervised learning pipeline holds huge potential by addressing computational bottlenecks and revealing new insights into complex data structures.

Diffusion (probability) models can be seen as a form of supervised learning. These models, specifically those used in generative tasks (like image generation), are based on a process where data is gradually corrupted with noise in a forward diffusion process, and then the model learns to reverse this noise process to recover the original data in the reverse diffusion process. Again, variants of HHL can provide quantum algorithms for diffusion models for which there is evidence for a quantum advantage.

3.1.2 - Reinforcement learning

Reinforcement learning (RL) is a branch of machine learning where an agent learns to make decisions by interacting with an environment to maximize cumulative rewards, outside the classification of supervised and unsupervised learning. It is used in many fields like robotics and large language models. Unlike supervised learning, which relies on labeled data, RL focuses on trial-and-error learning, using feedback from actions to improve future performance. A key challenge in RL lies in its computational complexity, as many RL problems are formulated using partially observable Markov decision processes (POMDPs), which are computationally demanding. This complexity makes finding optimal solutions computationally intractable for large-scale problems, further highlighting the need for efficient algorithms and approximations. Practically, RL relies on extensive datasets and lengthy training times, particularly in state-of-the-art models RL models which use deep neural networks. These hurdles underscore the potential of quantum computing to address the computational bottlenecks inherent in RL.

Efficient quantum training of classical models

One promising avenue involves using quantum computing to enhance the training of classical reinforcement learning models by employing hybrid quantum-classical architectures. Many RL applications, such as robotic manipulation and navigation tasks, rely on Actor-Critic frameworks, where the Critic evaluates actions to stabilize policy updates for the Actor. The Critic’s role is vital in these systems, as accurate value estimation improves learning stability and convergence rates.

However, achieving this in complex environments demands high computational resources and extended training times. Hybrid architectures address this by integrating quantum neural networks, which act as the Critic, with classical networks for the Actor. Quantum neural networks used as Critic are expected to better capture high-dimensional patterns, better trainability and improved stability, as well as better generalization during the Critic's evaluation. Once training is complete, the quantum Critic is removed, leaving a fully classical deployment system, without the need of a quantum computer during test. This design maximizes the practical utility of RL models while leveraging potential quantum advantages where they are most impactful during training.

In Actor-Critic architectures (classical) the critic is only used at time of training, but the neural network (NN) which outputs the action is the actor only. So, in this case there is no "transfer", because the quantum critic plays a role only in the training phase in both cases, classical and quantum.

Learning with quantum models

Another significant approach concerns quantum systems that replace the RL agent entirely, requiring quantum computation during both training and inferences. In this context, parameterized quantum circuits (PQCs) act as quantum agents, processing information and optimizing decisions in environments where classical agents struggle. These quantum agents are anticipated to be particularly advantageous in scenarios involving optimization in high-dimensional spaces such as those encountered in protein folding optimization, where the state space is massive and training a traditional RL agent is extremely challenging. They may also address problems with inherent quantum properties, such as quantum synthesis or compilation—optimizing the execution of a quantum algorithm on specific quantum hardware.

While this approach holds transformative potential, it also faces significant challenges. These include the current limitations of quantum hardware, issues with scalability, and the practicality of deploying quantum computers in contexts where the agent must be dynamic and actively interact with the environment. Nonetheless, quantum reinforcement learning agents represent a frontier in combining quantum computational power with RL's adaptive frameworks to solve problems beyond the reach of classical methods.

3.1.3 - Unsupervised learning

Quantum unsupervised learning seeks to leverage the principles of quantum computing to tackle the challenges of unsupervised learning tasks, which encompass clustering, dimensionality reduction, and generative modeling. One of the key distinctions within this field lies in the type of algorithms being developed: some aim to accelerate existing classical routines, while others introduce entirely new quantum-native methods. For instance, the q-means algorithm is a quantum alternative to the widely used k-means algorithm that provides a potential exponential speedup in running time while maintaining consistency with its classical counterpart.

Similarly, quantum algorithms for spectral clustering, a powerful technique for uncovering complex cluster structures, have been proposed, offering a theoretical polynomial speedup compared to the classical runtime. Despite the progress and promises surrounding quantum unsupervised learning, several challenges remain. One of the primary limitations is the difficulty of loading and processing high-dimensional classical data into a quantum state, particularly in the absence of ideal amplitude encoding that could be enabled someday by quantum memory models like quantum random access memory (qRAM). Additionally, while PQCs show promise, they face training challenges such as barren plateaus, where the optimization landscape becomes flat and hinders training. Currently, they remain impractical for addressing significant classical tasks.

Efficient quantum training of classical models

Quantum unsupervised learning utilizes quantum algorithms to identify patterns, clusters, or structures in unlabeled data. Clustering is a key methodology in unsupervised learning that groups data points into clusters based on their similarity. The goal is to partition data into groups where points within the same cluster are more similar to each other than to those in other clusters. Clustering typically works by identifying representative points, called centroids, for each cluster during training. Data points are iteratively assigned to the nearest centroid based on a predefined distance metric (e.g., Hamming, Euclidean).

Centroids are then recalculated to minimize the overall distance between points and their assigned centroids, repeating this process until convergence. Convergence occurs when cluster assignments or centroids stabilize, or when an objective function, such as the sum of squared distances, stops improving. This iterative process ensures that the algorithm identifies a stable grouping of data points that reflects the training data's structure. However, calculating these distances can be computationally expensive, particularly in large datasets with high number of features, as it may require evaluating all pairwise distances between data points.

Quantum clustering can address these challenges by leveraging the power of quantum algorithms, particularly in the computation of distances, which is at the core of clustering. One notable technique is the quantum swap test, which enables efficient calculation of the similarity (or distance) between quantum states. By employing the swap test, a quantum computer can determine the inner product between two data points encoded as quantum states, allowing potentially for significantly faster similarity comparisons compared to classical methods.

A prominent example of this is the acceleration of the k-means algorithm, where quantum methods can exponentially reduce the time complexity associated with calculating distances between data points and centroids by means of FTQC. Classical k-means involves iterative assignments of data points to clusters and recomputation of centroids, with each iteration requiring computationally expensive distance evaluations. A quantum variant performs these operations more efficiently. For instance, distance computations that scale linearly with the number of points and dimensions in classical approaches can scale logarithmically in quantum versions. Furthermore, the algorithm can encode cluster assignments in quantum states, enabling sampling and statistical analysis without directly revealing the entire dataset. This quantum advantage would make unsupervised learning feasible for high-dimensional, large-scale datasets, offering significant speedups and enhanced scalability in real-world clustering applications under the assumption that the cost of encoding classical data into quantum states is negligible

Importantly, while the calculation of the centroids can leverage quantum algorithms, one can also consider extracting this information from the quantum system and using a fully classical procedure to assign new points that were not originally part of the dataset. This method is effective only if the training set is sufficiently representative and the number of new points is limited, ensuring that the computational cost of distance calculations remains feasible for classical resources. By integrating quantum techniques into the training process, clustering models can benefit from significant speedups in distance computation, making quantum clustering a powerful approach for tackling large-scale and high-dimensional data.

Learning with quantum models

When adopting a quantum algorithm for both the training and testing phases of clustering, quantum methods can be leveraged to evaluate similarities, compute distances, and refine decision boundaries in high-dimensional spaces. This approach is particularly beneficial when the training set is limited, and the amount of unseen data is significantly larger. However, its feasibility depends on the scalability of quantum resources and the cost of quantum-classical transitions.

A more structured form of unsupervised learning, image segmentation, also encounters major computational challenges when relying on classical methods. Supervised segmentation relies on large amounts of labeled data, but obtaining high-quality annotations is costly, time-consuming, and often inconsistent due to human subjectivity. In many domains, such as medical imaging and remote sensing, ground truth labels are sparse, noisy, or unreliable, limiting the effectiveness of deep learning models trained on them. Unsupervised methods, including thresholding, clustering, region-growing, and edge detection, attempt to bypass the need for labels but come with their own limitations. These approaches often struggle with scalability, robustness, and adaptability to complex or high-dimensional data. More advanced graph-based methods, which represent images as graphs and segment them through optimization techniques like minimum cut or spectral clustering, offer a more structured approach but suffer from high computational complexity, particularly when solving large-scale combinatorial partitioning tasks.

Quantum computing presents a compelling alternative by leveraging quantum annealing to efficiently solve the combinatorial optimization problem underlying graph partitioning. This approach reformulates segmentation as a quadratic unconstrained binary optimization (QUBO) problem, allowing it to be executed on quantum annealers.

A key advantage is that in many cases, the grid structure of the graph representation of an image aligns well with the topology of quantum annealers, such as D-Wave's Pegasus architecture. This enables efficient embedding, allowing the solution of QUBO problems with thousands of variables while maintaining a highly favorable ratio between logical binary variables and physical qubits.

By encoding pixels as nodes and similarities or dissimilarities as edge weights, quantum annealing can rapidly find optimal partitions, outperforming, in some circumstances, classical solvers in runtime, particularly for large and complex datasets. This capability is especially useful in scenarios with limited or noisy labeled data, where traditional supervised learning struggles. Furthermore, while current implementations focus on quantum annealing, the same segmentation framework can be extended to gate-based quantum computing, leveraging qubit-efficient variational quantum algorithms for further advancements in large-scale segmentation tasks. However, while quantum methods show empirical advantages on specific scenarios, their extensive practical adoption depends on continued improvements in quantum hardware scalability, noise reduction, and efficient quantum-classical hybrid workflows.

3.2 - Research directions

3.2.1 - Learning models

Focus on quantum-process-based learning models: Most recent research in quantum machine learning (QML), particularly for near-term quantum systems, has centered on designing learning models around quantum processes. These include models such as quantum Boltzmann machines and parametrized quantum circuits used as distribution generators (e.g., quantum circuit Born machines) or as function approximators in supervised and reinforcement learning tasks. Considerable progress has been made in understanding the mathematical properties of these models, including their expressivity, complexity, and trainability bottlenecks. Recently, efforts have also focused on identifying the boundaries between trainable and non-dequantizable models, which contributes to a deeper understanding of when these quantum models can provide practical value. However, a critical question remains: why use these quantum models? The answer likely lies in future research, particularly with larger quantum systems where these models could demonstrate clear advantages.

Proving learning separations: In parallel, a smaller but impactful line of research has concentrated on proving separations between classical and quantum learning, focusing on specific tasks where QML provides provable advantages. Initially developed for cryptographic tasks, these separations have been generalized to more complex problems in quantum many-body systems, establishing strong connections between QML and quantum complexity theory. The concept of "data-hardness" emerges here, referring to tasks where the complexity of the underlying correlations makes the problem accessible to quantum learning but not classical learning, even with data. In other words, it can now be proven that are problem domains where QML has advantages, even though these cases remain limited and abstract, in that they usually apply to highly structured data in order to make cryptographic proof tools available. The broader range of machine learning tasks whose efficiency cannot be mathematically characterized - which is the case for most of machine learning (ML) in practice - remains an area awaiting larger quantum system, and more NISQ-friendly methods.

Research motivation and objectives: by combining these three key developments in QML, this project aims to identify new domains of application where quantum learning can achieve significant breakthroughs and in the near-term. Specifically, the project will focus on quantum methods for extracting useful features from hard-to-learn datasets, with applications in many-body physics and related computational tasks, in the cases they are likely "data-hard."

The ultimate goal is to develop quantum-assisted ML techniques that can outperform classical methods in certain well-defined domains of physics and beyond.

Short-term goals

- **Identify classes of features extracted in the quantum phase:** Develop an understanding of what makes certain features easier to extract using quantum machines.

- **Target domains in many-body physics:** Explore areas where quantum feature extraction could provide breakthroughs, such as:
    - Condensed-matter systems with highly correlated interactions and critical systems where classical ML methods fail.
    - Exotic phases of matter, such as systems with topological order or symmetry-protected phases.
    - High-energy physics applications, particularly in tasks related to quantum chromodynamics, where classical simulations are inadequate.
    - Quantum control problems, where quantum feedback mechanisms could provide superior optimization capabilities.

Mid-term goals

- **Transition from toy models to real-world applications:** Apply these methods to real-world problems in material sciences, solid-state physics, quantum chemistry, and high-energy physics (HEP). The aim is to move beyond theoretical models and demonstrate the utility of quantum-assisted learning in practical domains.

Long-term goals

- Optimizing material properties in quantum chemistry or solid-state physics.
- Anomaly detection: Analyzing HEP experiments data with hybrid quantum-classical models to extract novel insights or improve predictive accuracy.

3.2.2 - Quantum artificial intelligence for algorithmic discovery

Quantum information processing bears the promise of solving some hard computational problems or of improving their ability to perform distributed tasks. While concrete examples of such algorithms and protocols have fostered the development of the field during the past 30 years, designing new ones is both a necessity and a challenge. Because there is no a priori criterion that will ensure the existence of an efficient quantum solution for a given task, quantum algorithmic development is akin to searching a needle in a haystack.

The goal here is to leverage quantum and classical artificial intelligence to specifically address the issue of discovering and designing novel quantum algorithms and protocols. It involves utilizing quantum intelligent agents, specifically quantum neural networks, within a supervised and reinforcement learning framework. These agents will be trained to optimize quantum circuits that perform specific operations. Initial results have demonstrated that this approach produces optimal quantum circuits for tasks where quantum computers outperform classical ones such as the quantum Fourier transform (QFT), or nonlocal games (e.g. CHSH game). In addition, they took into account hardware constraints which suggest that these agents have not only the potential to uncover new, efficient quantum algorithms, a task that has proved to be a formidable challenge, but also to truly codesign.

Short-term goals

- Use the above quantum intelligence framework to re-discover known quantum algorithms and protocols (QFT, Grover's search algorithm, CHSH game, quantum key distribution, etc.).

Mid-term goals

- Discover novel quantum algorithms and protocols, for example, new algebraic transforms, cryptographic protocols, error correction codes, etc.
- Discover circuits that inherently provide noise-robustness via suitable constraints during training, in an application-agnostic fashion, in order to use such circuits in the NISQ era.

Longer-term goals

- Use AI to co-design quantum algorithms and quantum hardware, optimizing the necessary resources and bringing forward the quantum applications era.

3.2.3 - Quantum data pre-processing

Classical data processing methodologies refer to the traditional approaches used in computing for the preparation and analysis of data. These methodologies encompass a variety of tasks, including but not limited to, data cleaning, which involves the correction of errors or inaccuracies in the dataset. Another key task is feature selection, which aims to determine the most significant variables for the analysis at hand.

Additionally, normalization processes are applied to adjust the data's scale, enhancing the performance of models used for analysis. Techniques such as dimensionality reduction are also significant; they involve simplifying intricate datasets by reducing the total number of features while maintaining the dataset's critical information. These methods are instrumental in ensuring that data is both accurate and pertinent, rendering it suitable for further analytical processes. This makes them indispensable tools in the realm of machine learning and data science.

Quantum data preprocessing builds upon classical techniques by using quantum computing capabilities. This process entails the encoding of classical data into quantum states, with the potential goal of enabling more efficient storage and manipulation of information. Furthermore, it involves the development of methodologies capable of processing quantum states to extract meaningful information. Quantum algorithms could achieve computational advantages in handling large-scale and high-dimensional datasets, where classical methods become computationally expensive or infeasible. One fundamental aspect of quantum data preprocessing is quantum state preparation, where classical data is mapped onto a quantum system using techniques such as amplitude encoding, basis encoding, angle encoding, and so on. These representations could enable parallel processing and efficient transformations, reducing the computational cost of subsequent analyses.

Quantum algorithms, such as quantum principal component analysis (PCA), could allow for faster dimensionality reduction by using quantum linear algebra techniques. In particular, this approach offers an exponential speedup over classical PCA by leveraging quantum density matrix exponentiation. While classical PCA scales polynomially with system dimension, making it costly for large datasets, quantum PCA extracts principal components efficiently using multiple copies of a quantum system's density matrix. This approach is especially beneficial for low-rank matrices, as it identifies dominant eigenvectors without requiring full matrix construction—provided a fault-tolerant quantum computer is available.

Similarly, quantum feature selection methods, utilizing variational quantum circuits and optimization techniques, may allow us to identify relevant variables more efficiently than their classical counterparts. Indeed, by encoding data into quantum states and optimizing parameterized quantum gates, these methods strive to explore complex feature interactions more efficiently than classical algorithms. Quantum parallelism enhances correlation detection and the efficient evaluation of multiple feature subsets, potentially leading to faster and more accurate feature selection.

Quantum distance metrics and kernel methods also play a crucial role in the preprocessing phase of quantum machine learning by transforming data into a quantum-friendly format for further analysis. Quantum distance metrics, such as the Quantum Hamming Distance or Quantum Euclidean Distance, can compute the similarity between data points by exploiting quantum parallelism. This enables the identification of subtle correlations in high-dimensional data that might be difficult to detect classically. Similarly, quantum kernel methods map data into a higher-dimensional quantum Hilbert space, where complex relationships are more easily separated, enhancing the ability to perform clustering and classification tasks.

Finally, quantum generative models, including quantum Boltzmann machines and quantum generative adversarial networks, present novel frameworks for data augmentation, consequently positioning them as valuable instruments within data preprocessing workflows. The integration of quantum computing into data preprocessing pipelines has the potential to unlock new possibilities for the efficient handling of large, complex datasets of various application domains.

3.2.4 - Quantum optimization

Optimization is a fundamental problem in AI and beyond that involves identifying a point in a search space that minimizes or maximizes a given cost function. Many AI problems encounter extremely large search spaces, often due to combinatorial explosion in discrete domains, where the number of possible solutions grows exponentially. Additionally, in continuous search spaces, the lack of convexity in the cost function further complicates the process, leading to a potentially infinite number of solutions. This class of problems often is NP-hard in worst case complexity.

Quantum computing offers three promising avenues to address these challenges. First, quantum algorithms can encode and explore the search space of classical problems, potentially accelerating the discovery of optimal solutions, as envisioned with Grover's unstructured search algorithm in fault-tolerant quantum computing. For example, one can accelerate dynamic programming algorithms for the famous *travelling salesman problem* in this way, although actual speedups have to be assessed for practical use cases compared to best-in-class classical solutions. This usually leads to polynomial quantum speedups. Second, efficient optimization techniques are essential for training quantum circuits in variational algorithms, with the promise of enabling more effective solutions to optimization problems. Here, no proven separations are known yet.

Third, there are known problems where quantum algorithms can approximate instances of practically relevant problems such as integer linear programming well, while classical computer cannot approximate those instances well. In this sense, one can prove exponential separations of quantum over classical algorithms for sub-problems of hard optimization problems. Complementing this approach, decoded quantum interferometry connects combinatorial optimization problems, like sparse max-XORSAT, to decoding local density parity check codes, a task that can be efficiently solved using classical algorithms such as belief propagation. This approach leads to a quantum algorithm that surpasses classical optimization techniques, including simulated annealing. This section explores the intersection of quantum computing and AI in both directions, highlighting their synergies in tackling optimization challenges.

For natively quantum gradient descent algorithms: It is interesting to explore optimization algorithms finding approximate solutions efficiently that are natively suited to analog quantum simulators, as they can leverage the unique properties of quantum systems to address complex optimization challenges. For instance, quantum Hamiltonian descent (QHD) exemplifies this approach by mimicking gradient-based optimization through quantum dynamics. By utilizing quantum tunneling and other quantum mechanical effects, QHD enhances the ability to escape local minima and navigate challenging optimization landscapes. Such methods highlight the potential of quantum-native algorithms to complement or outperform classical techniques in optimization, particularly in scenarios involving highly non-convex or high-dimensional problems.

Evolutionary optimization: Evolutionary algorithms are search and optimization procedures inspired by the principles of natural selection and biological evolution. These algorithms emulate biological processes such as reproduction, mutation, crossover, and the survival of the fittest, thereby enabling the evolution of populations of candidate solutions over successive generations. In contrast to conventional optimization techniques, evolutionary algorithms operate on a population of solutions rather than a single point, thereby inherently facilitating parallel processing. This parallelism enables them to explore diverse regions of the search space concurrently, thereby markedly enhancing their capacity to evade local optima and identify global solutions. Therefore, quantum computing is an appropriate means of improving the efficacy of evolutionary algorithms by capitalizing on its inherent parallelism and computational capacity. Quantum computers can expedite pivotal operations within evolutionary algorithms, including fitness evaluations, mutation, and crossover procedures, by processing superpositions of potential solutions in a concurrent manner.

Quantum automated planning and scheduling (QPS) explores the integration of quantum computing into AI planning and scheduling tasks, focusing on quantum-supported planning (QP) and scheduling (QS). QP addresses tasks such as online planning, which involves real-time feedback, and offline planning, which is detached from execution. Under uncertainty, challenges such as partially observable Markov decision processes (POMDPs) arise, requiring complex strategies like conditional plans or utility-maximizing policies. These problems are computationally expensive, with tasks often being PSPACE-complete or worse. Initial quantum methods, such as quantum POMDP models and quantum Markov decision processes (QMDPs), propose theoretical frameworks but lack concrete experimental validation.

Similarly, QS tackles problems like job shop scheduling and its flexible variants (flexible job shop scheduling), which aim to optimize job assignments on multi-purpose machines to minimize objectives like production makespan. Quantum approaches, including quantum genetic algorithms, quantum particle swarm optimization, and hybrid MILP-QUBO (mixed-integer linear programming with quadratic unconstrained binary optimization problem formulation) methods, demonstrate potential speedups in specific instances, such as solving scheduling tasks for up to few hundred machines and jobs. Additionally, problems like bin packing have been addressed using quantum annealing, though current hardware limitations constrain their effectiveness. While early results highlight quantum advantages under specific conditions, further research is needed to identify scenarios where quantum methods outperform classical approaches.

Short-term goals

- Design and implement quantum-native optimization techniques and benchmark their performance against classical methods in AI applications such as combinatorial optimization and machine learning.
- Conduct medium-scale experiments to validate quantum-assisted scheduling and planning methods, identifying practical cases where quantum approaches offer computational advantages over classical techniques.

Medium-term goals

- Develop hybrid quantum-classical optimization frameworks and integrate them into industrial AI applications, such as logistics, supply chain management, and financial modeling, leveraging near-term quantum devices.
- Design tailored quantum evolutionary algorithms optimized for specific quantum hardware and benchmark their performance on real-world optimization tasks, such as scheduling, bin packing, or combinatorial design, demonstrating practical advantages over classical approaches.

Long-term goals

- Demonstrate a provable quantum advantage in practically relevant optimization tasks related to AI, such as large-scale automated planning, complex scheduling utilizing fault-tolerant quantum computers.
- Implement and validate quantum optimization techniques on large-scale industrial problems, using fault-tolerant quantum computers or specialized quantum hardware.

3.2.5 - Quantum reasoning

Reasoning is the process by which intelligent agents draw conclusions, make decisions, and solve problems based on available knowledge, rules, and observations. It allows agents to infer new information from existing facts, resolve uncertainties, and adapt to changing conditions. Reasoning can be classified into several types, including deductive reasoning (deriving logically certain conclusions from general rules), inductive reasoning (inferring general principles from specific instances), and abductive reasoning (finding the most plausible explanation for given observations). Intelligent agents implement reasoning through formal logic systems, such as propositional logic, first-order logic, and probabilistic logic, to construct structured representations of their surroundings. They encode entities, relationships, and evolving conditions using knowledge representation techniques like ontologies, semantic networks, rule-based systems, and Bayesian networks.

These structured models enable agents to process complex information, infer new knowledge, and update their beliefs dynamically. Intelligent agents can analyze representations using inference mechanisms, which include logical inference engines, probabilistic reasoning models, and machine learning-based predictors. These mechanisms allow agents to generate new insights, optimal decisions, and adaptive action plans. As their environments evolve, intelligent agents continuously refine their models to maintain accuracy, coherence, and responsiveness in real time.

Quantum computing could enhance these capabilities by taking advantage of quantum parallelism to encode and process knowledge structures more efficiently. It could enable agents to evaluate multiple reasoning paths simultaneously, thereby significantly accelerating the processes of inference and decision-making. Entanglement could enhance the representation of complex dependencies, improving context-aware reasoning and multi-agent coordination. Quantum interference has been shown to refine solutions by amplifying correct inferences and reducing computational errors.

The aforementioned quantum properties provide a fundamental advantage in fuzzy logic systems, probabilistic knowledge bases, and combinatorial optimization, thereby rendering AI agents more efficient in handling uncertainty, large-scale reasoning tasks, and highly dynamic environments. For example, in rule-based systems quantum computing could provide the ability to evaluate multiple rules and potential conclusions simultaneously using quantum superposition. Unlike classical systems, which process rules sequentially or in parallel with significant computational overhead, a quantum circuit can encode an entire rule set and explore all possible inferences at once. This could drastically improve efficiency in expert systems, automated reasoning, and legal decision-making, where complex rule dependencies must be evaluated rapidly.

Additionally, quantum entanglement can enable richer representations of logical relationships, allowing for more nuanced and context-aware reasoning. Another goal of quantum computing in reasoning could be related to probabilistic inference and Bayesian networks. Indeed, using quantum phase estimation and amplitude amplification, quantum algorithms can sample probability distributions potentially in a more efficient way than classical Monte Carlo methods. Moreover, other potential applications of quantum computing in reasoning are related to quantum-native reasoning architectures, where inference mechanisms exploit quantum entanglement and superposition rather than just mimicking classical logic. This could enable more advanced AI models capable of reasoning with fewer data and better than classical systems.

Short-term goals

- Develop quantum-assisted reasoning models to accelerate specific inference tasks (e.g., rule evaluation, probabilistic inference).
- Explore quantum-enhanced probabilistic reasoning techniques, such as quantum-assisted Bayesian networks and fuzzy logic systems.
- Implement proof-of-concept quantum circuits for evaluating multiple logical rules in parallel.
- Investigate hybrid quantum-classical approaches for automated reasoning and decision-making.
- Benchmark quantum algorithms against classical methods in expert systems and legal decision-making.

Medium-term goals

- Optimize quantum reasoning architectures for scalability and efficiency in large-scale reasoning tasks.
- Integrate quantum-enhanced inference mechanisms into AI-driven decision-support systems.
- Develop quantum-inspired techniques for multi-agent reasoning and context-aware decision-making.
- Improve quantum algorithms for probabilistic inference, using phase estimation and amplitude amplification.
- Establish practical applications of quantum computing in real-world reasoning tasks, such as legal analysis, financial modeling, and medical diagnosis.

Long-term goals

- Design fully quantum-native reasoning architectures that exploit entanglement and superposition for advanced AI cognition.
- Achieve quantum advantage in complex reasoning tasks, surpassing classical AI in efficiency and accuracy.
- Develop general-purpose quantum reasoning frameworks applicable across diverse AI domains.

- Explore quantum-enhanced learning models capable of reasoning with minimal data and handling extreme uncertainty.
- Integrate quantum reasoning with broader AI ecosystems, enabling next-generation autonomous systems with superior adaptability and intelligence.

3.2.6 - Quantum algorithms for multi-agent systems

Quantum multi-agent systems (QMAS) research explores the integration of quantum computing into the framework of autonomous agents and multi-agent systems. This integration involves primarily designing quantum-enhanced methods for coordination and cooperation among agents in centralized and distributed environments. These systems consist of multiple agents, either homogeneous or heterogeneous, that collaborate or compete in complex environments to achieve both individual and joint goals simultaneously. Despite the maturity of classical multi-agent frameworks, adapting these systems to leverage quantum computing is a relatively nascent field. Quantum-supported coordination methods have been proposed to enhance agent collaboration. Notable advances include quantum coalition protocols and contract net systems, which are methods used to negotiate and form agent coalitions in competitive or cooperative settings. Quantum coalition protocols have demonstrated reduced communication overhead and computational benefits over their classical counterparts. For example, quantum versions of coalition negotiation and resource allocation methods, such as quantum contract net protocols, offer enhanced privacy for agents while maintaining comparable computational efficiency. However, these advantages are solely theoretical and require further exploration.

Beyond theoretical constructs, with the advent of near-term quantum technology, QMAS have begun to address real-world problems on existing quantum hardware. However, many proposed methods have been tested only on simplified models, and their scalability and utility for complex, real-world problems require further investigation. Nevertheless, the ongoing research highlights the potential of quantum computing to address computationally intensive challenges in multi-agent systems, offering a glimpse into the future of AI in quantum-enabled environments.

One key area where QMAS show promise is in enhancing cooperation and coordination among agents, particularly in complex optimization tasks such as resource allocation, network optimization, and social network analysis. A fundamental challenge in multi-agent system (MAS) is coalition formation, where agents are grouped into teams or coalitions to maximize collective utility. Given the combinatorial nature of the coalition structure generation problem, this can be reformulated as QUBO problem, as demonstrated by BILP-Q algorithm and allows the use of quantum routines like the quantum approximate optimization algorithm (QAOA) and quantum annealing to explore large solution spaces presumably more efficiently. However, due to the current limitations of quantum hardware, scaling these approaches to larger problems remains an open challenge.

Hybrid quantum-classical methodologies address these limitations by optimizing the interaction between quantum and classical resources. For instance, GCS-Q employs a top-down approach where the grand coalition is iteratively split into smaller coalitions using quantum annealing to solve the computationally expensive bipartitioning problem. At each step, the algorithm formulates the bipartitioning as a QUBO problem and solves it on quantum hardware, while the classical part manages the overall coalition structure and ensures efficient recursion. Similarly, QuACS adopts a hybrid approach using QAOA to find the optimal bipartitions for induced subgraph games, combining quantum procedures in solving subproblems with classical logic to maintain scalability.

These hybrid methodologies delegate only the most quantum-suited tasks, such as solving graph-cutting problems, to quantum hardware, while classical components handle tasks that quantum systems cannot yet scale to. This strategic delegation results in a significant reduction in runtime and improved solution quality. Moreover, hybrid designs ensure that the advantages of quantum acceleration are leveraged without overburdening the limited qubit capacity of current hardware.

Future research should focus on qubit-efficient variational quantum algorithms, which further optimize quantum resource usage. These methods aim to minimize qubit requirements while maintaining quantum advantages in solving high-dimensional, combinatorial problems like coalition formation. By carefully integrating quantum and classical elements, hybrid quantum-classical approaches present a scalable path for addressing the computational challenges of coalition formation in MAS.

Short-term goals

- Develop qubit-efficient variational quantum algorithms to enhance coordination and decision-making in multi-agent systems (MAS) while mitigating current hardware constraints.
- Experimentally validate quantum-enhanced multi-agent strategies, such as coordination and negotiation protocols, using near-term quantum devices and hybrid quantum-classical methods.

Mid-term goals

- Collaborate with hardware developers to improve quantum architectures tailored for MAS tasks, focusing on optimizing quantum circuits for agent interactions.
- Design more efficient integration strategies that balance quantum and classical computation to expand the scalability of QMAS solutions in real-world AI-driven environments.

Long-term goals

- Utilize fault-tolerant quantum computing to enable entirely quantum-driven agent interactions, decision-making, and optimization, surpassing classical limitations in MAS applications.
- Deploy QMAS methodologies in high-impact domains such as autonomous systems, distributed AI, and dynamic resource management, leveraging advanced quantum hardware with improved stability and scalability.

3.3 - Use cases and applications

3.3.1 - Healthcare and life sciences

As of 2024, the fields of medicine and life sciences face several significant roadblocks that hinder the full potential of medical advancements. One of the primary challenges is the high cost and complexity of developing new treatments, particularly in areas like gene therapy and personalized medicine. While gene therapy offers incredible promise for curing genetic disorders, issues such as targeting the right cells and controlling treatment dosage remain significant obstacles that need to be addressed before these therapies can become more widespread.

Another major roadblock is the integration of AI and machine learning into diagnostics and treatment planning. While AI has made significant strides in improving diagnostic accuracy, the technology still faces limitations in data quality and interpretability. Ensuring that AI systems can provide reliable, unbiased, and explainable results is crucial for their broader adoption in clinical settings.

By using ab-initio simulations generated on quantum computers as training datasets for AI models, we could explore vast regions of chemical compound space that go beyond traditional bio-like molecules. This synergy between quantum computing and ML could significantly enhance drug discovery by providing a more comprehensive exploration of chemical spaces, identifying novel compounds, and accelerating the design process.

Quantum computing for accurate chemical simulations: Quantum computers could accurately simulate molecular properties at the quantum level, which is critical in drug design where small changes in molecular structure can drastically affect a drug's efficacy. Classical computers struggle with simulating the behavior of large, complex molecules due to the exponential increase in computational demands, but quantum computers may handle these tasks efficiently, allowing for precise ab-initio calculations.

Machine learning for efficient exploration: While quantum computing may generate highly accurate datasets, it is currently limited in scalability. Machine learning models, on the other hand, excel in handling large datasets and exploring vast solution spaces efficiently. By training ML algorithms on quantum-generated datasets, we can develop models that generalize well to new, unexplored regions of the chemical compound space, predicting the properties of molecules that lie far from the bio-like compounds traditionally studied in drug design.

Expanding the search space: Traditional drug discovery is often limited to molecules that are structurally similar to known bio-like compounds, which limit innovation. By leveraging ML models trained on quantum-derived data, we could explore chemical spaces that are far removed from this limited set, potentially uncovering entirely new classes of molecules with unique biological properties. This could lead to the discovery of novel therapeutics with mechanisms of action previously unknown to science.

Improving drug design: Quantum simulations may provide insights into complex phenomena such as protein-ligand interactions, electronic structure, and reaction mechanisms with high precision. When these simulations are used as training data for ML algorithms, the resulting models can predict chemical properties like binding affinity, solubility, and reactivity more accurately. This reduces the time and resources needed for experimental testing and accelerates the identification of promising drug candidates.

3.3.2 - Industry

Image analysis: Current state-of-the-art deep learning models for image analysis heavily rely on the availability of large and labeled datasets. However, with the rapid pace of data generation in fields such as medical imaging, autonomous driving, and satellite imagery, it is often unfeasible to label all available data due to the high costs, time, and expertise required for manual annotation. While unsupervised classical approaches, such as graph-based segmentation, provide alternatives for image processing without labeled data, they remain limited in practice because of the computational intensity required for handling high-resolution, real-world images. This challenge is especially evident in image analysis tasks like segmentation and motion detection, where processing extensive datasets, complex patterns, and inherent noise or inconsistencies further complicates the computational workload. Quantum algorithms present a compelling solution, with the potential to reduce computational loads and increase efficiency in tasks like graph-cut optimization and classification. Particularly in applications with noisy labels or where rapid processing is essential, quantum methods offer promising enhancements, addressing some of the primary limitations of traditional approaches in image processing.

Safe navigation in autonomous driving: The problem of collision-free navigation (CFN) for self-driving cars is a complex optimization problem, often modeled as a POMDP. Current state-of-the-art solutions rely on deep reinforcement learning, which, while effective, requires substantial computing resources and training time. QRL has emerged as a potential solution, showing faster convergence and improved stability in simplified environments. QRL methods leverage quantum computation, specifically PQCs, which have demonstrated polynomial improvement in parameter space complexity compared to classical deep Q-networks. However, current QRL methods have not been tested in complex real-world environments like CFN. The introduction of quantum components, specifically for the critic in an actor-critic framework, offers potential advantages in training complex RL architectures like Nav-Q, aiming to enhance trainability and stability without requiring onboard quantum hardware during testing.

Time series analysis: Time series data, prevalent in domains like finance, healthcare, meteorology, and industrial monitoring, presents unique challenges for AI due to its temporal dependencies, high-dimensionality, and potential for irregular sampling or noise. Classical methods, while powerful, often struggle to capture long-term dependencies or efficiently process massive datasets with intricate temporal patterns. Quantum methods, with their ability to process complex correlations and dynamics, offer promising advancements in time series analysis. Quantum algorithms could enable more efficient modeling of temporal patterns, possibly using less memory than their classical counterpart, faster anomaly detection, and enhanced forecasting accuracy by leveraging the native ability of quantum systems to handle high-dimensional data and optimize over non-linear relationships. Additionally, quantum-enhanced versions of classical techniques, such as recurrent neural networks or transformers, might provide improvements in processing and predicting time-dependent data across a range of critical applications.

Bin packing: Planning and scheduling problems are central to a wide range of industries, from logistics to manufacturing, and their complexity poses significant challenges as problem size grows. In logistics and supply chain management, the bin packing problem is a prime example. This task involves efficiently packing items into the minimum number of containers without exceeding capacity limits, a critical issue in optimizing warehouse storage, transportation, and data center resource allocation. Quantum algorithms, particularly QUBO-based techniques, hold promize for solving these problems more effectively than classical methods, potentially improving space utilization and operational efficiency.

Job shop scheduling: In industrial production, the flexible job shop scheduling problem is another key application. This task requires assigning jobs to machines to minimize production time, a complex challenge due to the vast number of potential job-machine combinations. Hybrid quantum-classical methods offer a practical approach by dividing the scheduling process into smaller, more manageable subtasks. These methods can enhance machine allocation, reduce production makespan, and adapt to dynamic industrial environments, making them valuable for advanced manufacturing systems in Industry 4.0. These applications demonstrate the potential of quantum technologies to revolutionize planning and scheduling in real-world settings.

Peer-to-peer energy trading: Peer-to-peer energy trading is a complex problem typically framed using multi-agent systems, where the agents are individual energy producers (e.g., households with solar panels) and consumers negotiating energy exchanges. Each agent's goal is to maximize utility, such as minimizing energy costs for consumers or maximizing revenue for producers, while ensuring the overall system remains balanced. The complexity arises from balancing supply and demand in real-time despite the intermittent nature of renewable energy sources, designing fair and efficient pricing mechanisms, and optimizing infrastructure use, such as storage and grid connections. Additionally, regulatory and social considerations, such as compliance with policies and encouraging participation, add further challenges. Tailored quantum algorithms could offer advantages by optimizing dynamic multi-agent interactions, managing large-scale data for real-time decisions, and enhancing the efficiency of computationally intensive tasks like pricing and grid optimization.

Electric vehicle charging management: The management of electric vehicle charging systems also relies on multi-agent systems, where the agents are vehicles, charging stations, and grid operators interacting to optimize energy distribution. Each agent has its own objective: the vehicles aim to minimize charging costs and waiting times, charging stations aim to maximize throughput, and grid operators aim to ensure grid stability and prevent overloading during peak demand. The primary challenges include scheduling charging sessions efficiently to prevent congestion, minimizing overall operational costs, and optimizing the geographical distribution of energy supply. The variability in electric vehicles arrival times, charging requirements, and station availability adds further complexity. Tailored quantum algorithms could enable efficient scheduling and routing optimization, improve real-time energy distribution, and balance the competing objectives of all agents while maintaining grid stability.

Dynamic environments in mobility and robotics: a multi-agent system framework is crucial for coordinating interactions between diverse entities such as vehicles, drones, robots, and their environments. Each agent, whether an autonomous car or a robot, must optimize its decisions in real time to achieve goals like collision avoidance, efficient path planning, or task completion. Challenges arise from the dynamic and uncertain nature of environments, requiring agents to adapt to changes in traffic, obstacles, or tasks. Quantum algorithms for multi-agent systems could offer significant advantages by enabling faster optimization of complex interactions, improving decision-making under uncertainty, and scaling efficiently with the number of agents. Applications include optimizing coordinated fleet movements, enhancing task allocations in robotic swarms, and accelerating real-time computations for dynamic and unpredictable environments, paving the way for more robust and efficient autonomous systems.

3.3.3 - Quantum physics

The rapidly evolving field of QML offers opportunities for tackling complex problems in physics, particularly in systems where classical ML techniques struggle, but also in other computationally hard tasks. Recent advances in QML highlight three key research directions, which collectively point to new ways in which quantum computers could outperform classical systems for specific learning tasks, especially in complex quantum systems. Here we aim to build on these developments and explore the potential for quantum computers to extract valuable insights from challenging datasets in well-defined domains of physics, including many-body physics, condensed matter, and high-energy physics.

4 - AI for quantum

Artificial intelligence, machine learning, and computational methods inspired by them are transforming mathematical and computational modeling techniques across the board, offering in many cases a novel perspective on what it means to create a model. It accelerates work flows and moves the needle on the question what intellectual task can be automated. Diligent use of these methods makes human knowledge work more meaningful in that it guides and supervises the AI.

For quantum technologies, AI is a key enabler with a potential that goes far beyond matching two current technologies. First, the probabilistic character of quantum physics can be matched to that of AI. Second, constructions in quantum technologies often defy intuition yet are based on mathematical equations. Third, the strong precision requirement of quantum technology requires testing against complex, multi-parameter models taking the tiniest disturbances into account. Finally, during its buildup, shortcuts to building up a quantum workforce are more than welcome.

Short-term, mid-term and long-term goals

- Enhancement of the control software stack with AI keeping quantum technologies at their maximum performance most of the time.
- Automated post-processing of data from quantum experiments to output useful data for humans.
- Development of automated design tools for quantum experiment realizations with low need for human correction.
- Automated quantum error correction with AI enhancement, improving both threshold and overhead.

4.1 - Discovery and optimization of quantum experiments

The routine operation of quantum computers will require new approaches for device design, optimal control, readout, efficient compilation, detection of noise mechanisms, and error correction. One example of enhanced compilation is ML models that can autonomously learn generic strategies to compress quantum circuits.

On a larger scale, variational quantum algorithms or other forms of quantum machine learning have been identified as a promising route to discover new algorithms in order to supplement the so far rather small number of algorithms with potential quantum advantage.

This illustrates that the machine learning toolbox provides manifold ingredients that can boost the hybrid quantum-classical operation for these purposes. In addition, ML is one of the most prominent tools for an efficient interface between classical and quantum systems both for experiment discovery and control as well as for the translation of classical data into a quantum state and the transpilation of classical into quantum algorithms. ML algorithms operating in hybrid classical-quantum hardware will also allow us to harness properties of quantum systems to devise energy-efficient control tasks.

Concrete examples of ML applications:

- Cross-architecture optimization of quantum devices and quantum experiments.
- Hybrid quantum-classical devices and control protocols.

Short-term goals

- Development of generative and reinforcement learning methods able to compile sequences of quantum operations, e.g. quantum circuits, given some input tasks, optimize the number of gates, and more generally, improve them based on hardware constraints.
- Creation of new experimental setups able to create target quantum states, an, in particular, interesting ones like GHZ and other entangled states, with fewer physical components or new platforms.

Mid-term goals

- Discovery of more efficient implementations of known quantum algorithms as well as better ansatz for e.g. quantum chemisty and quantum simulation problems.

Long-term goals

- End-to-end generation of quantum routines for input tasks, from the devising new quantum algorithms / variational approaches to its efficient implementation in the native gates of the given hardware.

4.2 - Simulation of quantum systems

The simulation of quantum many-particle systems on classical computers is one of the greatest challenges in physics, but extremely important for advances in developing functional materials and chemical processes. It has been demonstrated for a wide range of traditional ab initio methods how machine learning components can be incorporated for their advancement. One example is the encoding of the quantum wave function in the form of artificial neural networks (neural quantum states), which was found to efficiently capture complex quantum states.

These methods are currently pushing the state of the art and may be applied to outstanding problems such as high-temperature superconductivity and chemical reaction dynamics. By combining domain-specific knowledge of quantum mechanics with ML techniques, physics-aware ML can improve the accuracy and interpretability of predictions and models for quantum systems. One example of this is the learning of the hidden disorder landscape of a quantum device. Developing explainable AI for quantum physics could foster the adoption of quantum machine learning techniques as the paradigm for digital twinning of quantum systems, as well as identify new instances of quantum systems, which are out of reach to simulate with classical ML resources.

Concrete examples of ML applications

- Neural quantum states.
- Automated tailoring of properties of quantum materials.
- Digital twins for quantum systems.
- Classical shadows.
- Classical control of Quantum circuits.

4.3 - Analysis of quantum data

Data from quantum systems is particularly complex and hard to analyze. For instance, in quantum many-body systems, such as correlated electrons in solid-state systems, the relevant order at the basis of their functionality is often hard to extract. Recently it was found that machine learning tools can guide and extend such analysis by an unbiased exploration of all information. For example, artificial neural networks can be used to identify the relevant correlations from snapshots of quantum many-body systems on a lattice. This can result in the unbiased identification of the essential observables, which is particularly important in noisy settings or for exotic order, where traditional methods fail. In turn, generative modelling can provide crucial support in enriching data from quantum systems where extracting measurements is hardly accessible experimentally or is particularly time consuming. For that purpose, interpretable and explainable machine learning models will be of particular use to facilitate the discovery and understanding of such new observables.

Concrete examples of ML applications

- Discovering correlations and symmetries in quantum experiments.
- Physics insights through explainable AI and AI-assisted discovery.

Short-term goals

- Develop supervised and unsupervised learning methods able to characterize data arising from various types of quantum hardware.
- Introduce interpretable and explainable models and architectures able to re-discover known quantum theories and phenomena directly from data.

Mid-term goals

- Build autonomous pipelines able to create, directly from data, new theories and efficient descriptions of the data arising from quantum experiments.

Long-term goals

- Combine these methods with those proposing and controlling quantum systems to create better experiments from which to acquire high-quality from which to acquire new theoretical insights.

4.4 - Automated control and calibration of quantum technologies

Both quantum technologies and fundamental experiments rely on the precise manipulation and fabrication of a large number of quantum degrees of freedom. Common challenges are the efficient characterization, avoidance of inevitable hardware imperfections, and optimization of control strategies. As the qubit numbers grow and the applications become manifold, automation becomes essential and machine learning tools have been shown to be powerful for this purpose.

Examples include the fully automated tuning of quantum-dot devices, the automated optimization of entangling operations on a superconducting quantum processor could substantially improve their quality, which is crucial for near-term applications. ML will accelerate the process of quantum state preparation, gate operations, and measurement, leading to faster and more efficient quantum computation or quantum communication protocols on given hardware. Thus, the ML toolbox can be employed to extend the applicability of quantum devices to problems with many noisy parameters such as imaging, radar, or gravitational wave detection. ML methods can readily be utilised to characterize noise sources.

By facilitating the exploration of large parameter spaces, ML can also be used to design optimized experimental setups for specific quantum tasks, such as quantum communication or quantum computing.

Concrete examples of ML applications

- Efficient characterization, tuning, design, and control of quantum experiments.
- Mitigation and harnessing noise in quantum systems.
- Cross-architecture optimization of quantum devices; cross-platform certification.

Short-term goals

- Online control of quantum experiments via reinforcement learning and other methods, improving the stability and coherence of quantum systems.
- Development of better state preparations, operations and measurement patterns, enhancing the quality of current experiments.

Mid-term goals

- Development of fully-automated quantum laboratories, easing the preparation of multi-component experiments (e.g. calibration of optical tables and similar), for a wide variety of given tasks (i.e. multi-task ML agents). This would be combined with large language models (LLM) for easy user interaction.

Long-term goals

- It would relate to the previous but adding an extra step of physics discovery with ML models to know which experiment we really want to do.

4.5 - Trustful, robust, interpretable, and explainable AI for quantum technologies

Key considerations in applying machine learning to quantum science and technology are trust, robustness, interpretability, and explainability. While neural networks (NNs) have shown their power in various applications, their lack of transparency hinders the safe and reliable application of these algorithms to valuable quantum systems. Striking a delicate balance between leveraging advanced algorithms and mitigating risks is crucial for instilling confidence in automated control systems.

Additionally, interpretability and explainability of ML in quantum science are critical for uncovering decision mechanisms used by NNs when addressing complex quantum problems. To drive scientific discovery, it is vital to not only comprehend the outputs generated by ML algorithms, but to also understand the underlying principles and concepts that guide their reasoning. Understanding the factors contributing to a model's predictions allows scientists to assess reliability, validate solutions, and identify biases and errors in training data. Ultimately, this improves the robustness of quantum simulations and predictions. Further, extracting human-understandable knowledge from ML models is pivotal for driving breakthroughs in quantum science and technology. Efficient implementation of approaches that effectively contribute to validating findings, uncovering novel insights, and advancing quantum science through new discoveries is a significant challenge.

Application of AI onto quantum technologies is an excellent test scenario for these techniques, as AI discoveries can be retroactively explained, interpreted, and tested. In fact, first steps have been taken towards introducing genuinely quantum methods for providing tools for explainable quantum machine learning, based on *Taylor*-∞ that is a black-box explanation that exploits the Fourier picture of parametrized quantum circuits and quantum layerwise relevance propagation. Such tools can provide guidance on what it is that the quantum machine learns and provide valuable diagnostics for interpreting and explaining quantum models in the machine learning context.

Short-term goals

- Development of trustful AI models for the control of critical quantum systems.
- In line to unsupervised learning for the analysis of quantum data is the development of scalable explainable and interpretable models with specific inductive biases towards quantum data.

4.6 - Quantum error mitigation via post-processing by AI

Quantum algorithms are known to typically output probability distributions over possible measurement outcomes. These distributions encode the solutions to computational problems, such as optimization, simulation, or factorization. However, due to inherent hardware limitations, noise, and decoherence, the observed probability distributions often deviate significantly from their ideal counterparts. This discrepancy poses a major challenge for practical quantum computation, particularly in the current era of NISQ devices.

From a computational perspective, errors in quantum probability distributions emerge due to several factors. These include gate imperfections, state preparation and measurement errors, and decoherence caused by the environment. These errors distort the statistical properties of quantum algorithms, resulting in unreliable outputs. Unlike classical errors, which can often be corrected deterministically through redundancy, quantum errors require sophisticated mitigation techniques due to the no-cloning theorem and the fragility of quantum states. There are also worst case bounds known that limit the applicability of quantum error mitigation methods to basically log-log-depth quantum circuits in worst case complexity.

Recent advancements in artificial AI techniques, particularly machine learning, reasoning, and optimization, have emerged as powerful tools to correct and refine quantum probability distributions post-processing. Machine learning models, such as neural networks and Gaussian processes, have the capacity to learn error patterns from experimental quantum data and apply corrections to infer the ideal distribution. Furthermore, reinforcement learning and Bayesian reasoning can enhance error mitigation by dynamically adapting correction strategies based on observed system behavior.

Approximate reasoning, as exemplified by fuzzy logic, plays a pivotal role in the mitigation of quantum errors by addressing the inherent uncertainty and imprecision in the identification of quantum error patterns. By representing quantum errors in terms of multi-value logic, it becomes possible to identify, model, and correct errors in a more flexible and robust manner. To illustrate this point, fuzzy inference systems or algorithms for approximate clustering can be utilized to dynamically adjust post-processing parameters based on the degree of noise present in quantum outputs. This, in turn, improves the reliability of error mitigation techniques.

Furthermore, optimization methods enable the precise adjustment of post-processing parameters to restore fidelity to the quantum outputs. For instance, evolutionary optimization, incorporating genetic algorithms, differential evolution and memetic algorithms, offers a robust approach to refining quantum error mitigation strategies. These techniques can optimize the parameters of mitigation operators, machine learning models, fine-tune fuzzy logic-based correction mechanisms, and identify novel error mitigation strategies by simulating evolutionary processes. By repeatedly selecting the most effective mitigation techniques based on their performance, evolutionary optimization ensures that the correction methods continuously improve over time.

The utilization of artificial intelligence (AI)-driven techniques has emerged as a promising approach to enhance the reliability of near-term quantum computations. This enhancement is achieved without necessitating additional quantum resources, thus demonstrating the efficacy of AI-driven techniques in quantum error mitigation.

Short-term goals

- Develop machine learning models to learn and correct quantum error patterns from experimental data.
- Implement fuzzy logic-based approaches for identifying and mitigating quantum errors.
- Apply evolutionary optimization techniques to refine error mitigation strategies.
- Benchmark AI-driven post-processing methods against traditional quantum error mitigation techniques.
- Develop hybrid AI-quantum frameworks for improving the fidelity of quantum probability distributions.

Mid-term goals

- Develop AI-assisted dynamic parameter tuning for post-processing quantum outputs.
- Investigate scalable AI techniques to handle increasing quantum system sizes.
- Integrate AI-driven quantum error mitigation into practical NISQ applications, such as optimization and simulation.

Long-term goals

- Achieve AI-enhanced quantum error mitigation techniques that enable practical quantum advantage.
- Establish AI-quantum co-optimization strategies for next-generation quantum processors.

4.7 - Quantum error correction syndrome detection

A prominent task in large scale fault-tolerant quantum computing is the decoding of error syndromes into the most likely error pattern. This is a natural field of application of classical AI, which needs to be carefully benchmarked against more traditional algorithms such as minimum-weight matching. Moving the needle in error correction can have a dramatic impact in the reduction of error correction overhead. But we must ensure that machine learning based error syndrome detection scales well with large distance error correction codes.

Short-term goals

- AI-based quantum error decoders validated on digitally simulated experiments.

Mid-term goals

- AI-based quantum error decoders validated on experiments.

Long-term goals

- Practical fault tolerant quantum computing controlled by AI-based classical infrastructure.

4.8 - Quantum architecture search with machine learning for near-term algorithms

Variational quantum algorithms are widely used in the NISQ era to solve machine learning problems such as classification, prediction and generative tasks. Parameterized quantum circuits are crucial component of variational quantum algorithms (VQAs) and deep parameterized circuits encounter trainability issues, such as barren plateaus, quantum hardware constraints and noise further worsen the performance of variational quantum algorithms. To address these issues, there is need to design efficient circuits with optimal set of parameters tailored to underlying problems and quantum hardware, such a task of generating circuits for a specific problem also referred as quantum architecture search or automatic design of circuits.

In early 2000s, researchers used evolutionary algorithms to find efficient quantum circuits. Since then, various ML and optimization techniques such as deep reinforcement learning, Bayesian optimization, adaptive methods such as quantum autoencoders for data compression, ADAPT_VQE and other attention-based models are used to generate optimal circuits. Many quantum architecture search techniques are inspired by neural architecture search and can be used to generate an efficient ansatz which can help in tacking issues associated with VQAs, co-designing algorithms for specific quantum hardware or finding new efficient and better algorithms.

Short-term goals

- Use reinforcement learning and generative models for efficient circuit sampling.

Mid-term goals

- Generate optimal circuits and new quantum algorithms for industrial use-cases or applications.

Long-term goals

- Use AI techniques to discover new strategies and better quantum algorithms.

5 - Foundational questions

While much of the current interest in machine learning and AI stems from applied and practical considerations, fundamental questions about learning systems have in fact been a driving force of much of the progress. The importance of fundamental physics perspectives on traditional ML, which have been critical in achieving the current unprecedented points of accuracy, has been acknowledged even by 2024 Nobel Prize in physics. Statistical mechanics has led to the development of energy-based techniques and, then, diffusion models have achieved game-changing results in terms of data generation. Insights from studies on the renormalization group and mean field theory have been used in modeling ML/AI behavior.

Given the fundamental role that classical physics had in the development of classical AI and machine learning, it is tempting to ask whether novel quantum machine learning paradigms can be developed from foundational studies in quantum physics and related disciplines. For instance, quantum statistical mechanics and quantum thermodynamics may have a role as crucial as that of classical statistical mechanics for classical AI. The understanding of correlations in dynamical quantum many-body systems may inspire novel quantum algorithms and approximation methods. Progresses on foundational questions can also help defining novel approaches to deal and interact with "data" in the quantum world and to better understand learning in a fully quantum setting. Example questions include the different interpretations of quantum physics, the interplay between the quantum and the classical world, the role of the measurement postulate and all the limitations that come from that, such as the no-cloning theorem, but also the emergence of the wave-function collapse from other, more fundamental theories.

In this section we speculate on not fully understood, and yet-to-be discovered connections between physics, (quantum) machine learning and AI.

General goals

- Explore novel quantum machine learning paradigms inspired by foundational studies.
- Develop a fully quantum AI model, leveraging principles like superposition, entanglement, and non-locality to redefine learning in a quantum world.
- Better understand learning in the quantum realm, addressing fundamental questions such as the role of no-cloning theorem, quantum measurements and the wave-function collapse.
- Investigate quantum learning as a physical process, e.g. whether interacting quantum many-particle systems can naturally perform learning tasks through their natural dynamics.
- Apply physics principles to address safety, robustness, explainability, and interpretability in quantum AI systems, ensuring they are reliable and aligned with desired outcomes.
- Investigate quantum analogs of classical statistical-mechanical models (e.g., Hopfield networks) and their potential for storing and retrieving quantum information.
- Address challenges in the quantum agent-environment paradigm, such as entanglement between agents and environments, and develop frameworks for quantum generalizations of Markov decision processes.

5.1 - Physics and (quantum) machine learning

In this line of investigation, we reflect on the possible new connections between physics and learning systems. The question of new connections can be raised by considering the cutting edge and future challenges of even classical machine learning.

5.1.1 - Toward understanding complex and heterogeneous systems and toward general AI

Most progress in ML has been achieved by giving up on the dream of general AI, and by focusing on specific sub-tasks.

Recent developments in e.g. large language models foreshadow the return to the original challenges, where complicated learning systems (which are not just a large neural network, but a complex transformer architecture) fulfill ever more complicated roles. It is a question whether new ideas in physics modeling of complex systems can lead to new insights.

5.1.2 - Safe AI: robustness, explainability, interpretability, alignment and other

While the previous decades were hyper focused on performance in terms of accuracy, it is abundantly clear the next phases will require a more complicated metric. Arguably the most important features of AI we are interested in pertaining to safety of AI systems, in various contexts. The simplest cases include verification tasks (ensuring an AI system will never fail in some catastrophic way), robustness (stability to small random or adversarial perturbations of inputs), explainability (the capacity to reason about why certain decisions were made) and so on. It is a question whether physics principles and logic can be used to shed light on the limitations and perspectives of achieving AI systems which are safe.

5.1.3 - Quantum AI

Quantum machine learning and AI methods, being developed in the context of quantum information, are already naturally connected to aspects of quantum physics. However, it is unclear whether other parts of quantum theory, such as quantum thermodynamics, quantum statistical mechanics, condensed matter, ideas can be used to elucidate the learning processes of quantum mechanical systems. This remains another challenge.

5.2 - Machine learning and AI in a quantum world

Another fundamental question stems from the very definition of quantum machine learning. This field does not merely lie at the intersection of quantum information, quantum physics, machine learning, and artificial intelligence but rather somewhere in their union. This subsection attempts to address a speculative use case where everything can be quantum: the model, the training and inference algorithms, the data and possibly even the labels. Everything can be quantum in nature and all the data, after appropriate pre-processing, can be transformed to a non-trivial quantum state (e.g., by identifying high correlations through the incorporation of entanglement techniques). Therefore, it is crucial to foster the creation of methodologies and techniques that address quantum artificial intelligence at a level that is fully encompassed within quantum physics. Ultimately, machine learning can be understood as information processing. Quantum information can be very different from classical information. Genuine quantum principles and phenomena such as superposition, entanglement, non-locality, contextuality, and other quintessential facets of quantum information may alter both the definition and the meaning of learning in a fully quantum world.

The first class of foundational issues in QML arises when we consider various machine learning modalities (supervised, unsupervised, reinforcement), and allow data (inputs) and or labels (outputs) to be promoted to genuinely quantum states. The long-term objective here is to build a fully quantum AI model where all data, training algorithm and inference system are fully quantum.

We identify a number of classes of questions:

5.2.1 - How to define learning

For simplicity we will illustrate the questions on the cases of supervised learning and reinforcement learning. Early investigations into supervised quantum learning date back to the 2000s, where multiple copies of quantum states were considered as inputs. This setting has been explored extensively, yielding many interesting results. However, the exact similarities and differences between classical and quantum supervised learning remain unclear. For instance, it is still unknown what theoretical or practical limitations might exist when working with quantum data sets or when mapping classical supervised learning problems onto quantum analogs. If the learner uses up all the quantum data during the training phase, then the learning process is essentially classical, as the training set becomes a classical map.

Hybrid strategies where classical ML methods are paired with quantum measurements to extract information from quantum data may be problematic, at least in the worst case, due to exponential complexity of full tomographic methods. Recent advances in less general, but more efficient methods based on, e.g., classical shadows, quantum kernel methods, quantum Boltzmann machines or tensor networks, may be efficient for specific tasks, but the general principle is still missing. Learning strategies based on quantum memories are more likely to achieve probable speedups. Extra care must be taken if the final learned model is stored in a quantum memory, as the latter can neither be copied nor broadcasted. However, there are intriguing cases—such as gentle measurement techniques in shadow tomography—that suggest quantum states can retain usefulness even after partial measurement. Determining scenarios where supervised quantum learning algorithms must retain quantum training data for future tasks remains an open foundational problem. More fundamental problems emerge when we consider genuinely quantum label sets. This scenario introduces significant conceptual challenges.

5.2.2 - Learning as a physical process

Every quantum algorithm can be expressed as a suitably discretized evolution of a physical process, e.g. with time-dependent Hamiltonians or maps. A natural question is then whether we can define an interacting quantum many-particle system, e.g. with many qubits or other physical particles, whose natural evolution performs the learning task in a fully quantum world. Different “natural” evolutions are possible, e.g. via unitary dynamics or via an adiabatic evolution where the system approximately remains in the ground state of some evolving Hamiltonian. A possible example of this research line is about finding quantum analogies to the relationship between classical statistical-mechanical models and associative memories (Hopfield networks). Specifically, it is not clear if quantum channels (Hamiltonians, or open systems) can be similarly used to store and retrieve quantum information. Speculative examples may include storing/retrieving information in ground states of some complex quantum Hamiltonian (e.g. transfer Ising, where non-commutative terms enable a rich phase diagram), or in non-local correlations that are created in the scrambling dynamics of chaotic quantum systems.

5.2.3 - Learning agent-environment paradigm challenges

Reinforcement learning, and more generally the learning agent-environment paradigm also encounters numerous challenges when "quantized". In a fully quantum RL setting, both the environment and the agent may become entangled over the course of their interactions. This entanglement complicates the very notion of a “history” of interactions because measurements of this history collapse superpositions and potentially interfere destructively with the learning process. One partial solution is to redefine learning in these contexts using quantum generalizations of Markov decision processes (MDPs), as seen in works on quantum observable MDPs.

However, these frameworks typically assume quantum actions are represented as classical descriptions of quantum operations, rather than as quantum states themselves. A more general formalism—allowing agents to generate and act with quantum states—remains underexplored. Quantum processors, e.g. based on port-based teleportation, can be used to write the program in a quantum state, but no efficient training method is known to date. The question of information extracted during learning is also crucial. If an agent gathers information from a quantum environment, it must operate within the constraints of quantum mechanics—e.g., the no-cloning theorem. This raises fascinating foundational questions: can knowledge itself be defined in terms of quantum states that cannot be shared or copied? While this notion aligns with quantum mechanics, it challenges classical intuitions about knowledge transfer and collaboration in machine learning.

Ultimate questions here go in far, from the limits of quantum autonomous agents to learn new quantum physics, to in principle the influence learning may have on foundations of quantum mechanics, e.g. in the definition of an observer which performs a measurement.

6 - Building bridges between quantum and AI

We believe it is now the right time to invest in research at this emerging interface of quantum science and machine learning so that the EU can remain competitive with the US, Canada, and China in developing next-generation quantum technology. Patents for machine learning applications in quantum computing are already picking up speed, but mostly in the US. The quantum flagship has put Europe in a strong position and a broad funding initiative for machine learning in quantum science will enable Europe to take on the lead in these new developing technologies.

Funding needs to be both for fundamental and applied research projects in order to cover the full spectrum of developments. While the applications for optimal control are already being prepared for commercial exploitation by the first start-ups, ab initio computational methods are in a more exploratory phase, which requires funding of purely fundamental research for unleashing the full potential. The same holds for other applications such as the analysis of quantum data arising from experiments that might turn out seminal for the understanding of physics or for future technologies.

Facilitating the exchange between the ML and the quantum physics communities has the potential to transform both fields and interdisciplinary teams are needed to push beyond current boundaries. At the core of this proposed initiative is the merging of diverse communities, to bring together a heterogeneous range of views and ensure openness and diversity. For quantum science and technology to synergize with the field of ML and artificial intelligence, we shall need to bring together quantum experimentalists, quantum theorists, ML engineers, computer scientists, but also entrepreneurs and investors.

Open-source software, freely available and standardized benchmark data sets, model databases, and community challenges were central for the rapid advancement of machine learning techniques. Building on this experience, we believe that creating a similar ecosystem for machine learning in quantum science will likewise boost progress by removing barriers for interdisciplinary collaboration and optimally tapping the available potential. To this end, it is necessary to standardize quantum physics problems through interoperable and structured interfaces. Their role will be to enable sharing of experimental data and translating quantum physics problems into a common ML language. On the one hand, standardization will enhance the applicability of ML methods in both theoretical and experimental quantum physics, thus improving reusability, reproducibility, and comparability. On the other hand, the development of community-driven projects will create shared spaces, which provide interfaces as tutorials or documentations that help students and researchers to familiarize and strengthen cohesion between fields, and encourage interdisciplinary collaboration and cross fertilization.

Progress in this rapidly developing field requires the training of a next generation of researchers with expertise in quantum science and machine learning, e.g., via suitable doctoral networks. Additional training and educational resources, such as dedicated online platforms and training resources, and encouraging cross-field conferences and symposia, will simplify the access to state-of-the-art ML and further encourage its widespread adoption by quantum scientists. This can bridge the gap between theory and experiment, by facilitating a more seamless integration of theoretical modeling and experimental data analysis. Such educational programs at the interface of quantum science and machine learning will produce a workforce that is highly skilled in both forward-looking fields. This is not only fruitful for fundamental research, but also essential to keep replenishing industry with open-minded experts who transfer knowledge into competitive products and services.

Social media and science-communication strategies, as well as collaborations with creators, developers, and industry partners, will play a key role in making machine learning techniques in quantum physics beneficial to all of society. By placing engagement at the center of the research process we shall bridge boundaries between disciplines, facilitate the exchange of valuable knowledge with industry partners and policy makers, and improve the public perception of quantum science.

7 - Recommendations and challenges

Here are some of the key transversal operational scientific and technology challenges to address to advance the fields covered in this white paper.

7.1 - Theoretical work

Quantum artificial intelligence is still a relatively nascent domain. More theoretical insights are needed to better understand how to best employ quantum phenomenon to accelerate computing. It deals with determining theoretical and practical computing speedups, reduced requirement for training data and to obtain better results. Also, more theoretical work is needed to create efficient bridges between quantum and classical AI, like when a quantum computer is used to train a machine or deep learning model that is then run classically, like in embedded systems (car visions, etc). This work should also cover both fundamental and applied research and practical use cases, like in healthcare.

7.2 - Aligning with quantum hardware roadmaps

Implementing quantum-assisted artificial intelligence algorithms is highly dependent on the progress of quantum computing hardware. Particularly, it is related to the quality and quantity of available qubits. Quantum artificial intelligence capabilities will progress synchronously with hardware evolutions, such as advanced NISQ with better qubit fidelities, early fault-tolerant quantum computers (eFTQC) with about a hundred logical qubits, and beyond, with utility-scale FTQC quantum computers support thousand logical qubits. Algorithmic advances in quantum artificial intelligence will guide quantum hardware academics and industry vendors in adjusting their roadmaps. Likewise, quantum machine learning developers will synchronize their work with hardware vendors.

A second aspect deals with qRAM (quantum random access memory) research. qRAM will be an important enabling technology for quantum machine learning, particularly to address the pressing challenges with data preparation and loading.

7.3 - Estimating resources

In relation to the previous point, research work on quantum artificial intelligence as well as on the usage of classical machine learning for the development of quantum computing hardware and software must rely on careful resource analysis. It will help identify scaling issues and avoid massive energetic needs that society will struggle to provide. This transversal research will benefit from the involvement of the EuroHPC participating organizations.

As concern with the energy consumption of current AI and LLM solutions is growing, some academic and industry work should estimate, benchmark and optimize the energy consumption of both quantum artificial intelligence solutions and classical artificial intelligence tools used in quantum technologies.

When AI is used as a tool for calibration, error analysis for mitigation or correction, or other enabling tasks in quantum computing, it is crucial that its energy consumption does not negate the energy savings expected in quantum computing relative to classical HPC.

7.4 - Engaging classical AI specialists

Advancing the field of quantum machine learning as well as the use of machine learning in the context of quantum computing requires more engagement of the classical AI scientific community. Cross-discipline initiatives may be launched in education and community buildouts to encourage it.

This will also enable the quantum community to better qualify the challenges ahead with classical AI.

7.5 - Software engineering

The development of an EU-based known-how and competitive advantage on classical and quantum artificial intelligence should lead to the development of new software engineering tools. It deals with quantum code compilers and optimizers leveraging machine learning, tools enabling quantum code debugging and the likes. Proper usage of LLM-based software engineering will also be crucial for increasing the productivity of quantum software developers.

7.6 - Open science and industry competitiveness

The EU is highly challenged by US dominance in the AI field, at the hardware level (Nvidia) as well as with software and cloud infrastructures (OpenAI, Google, Meta, AWS). Meanwhile, a vibrant innovation ecosystem works well with open science and research processes. We have in mind that open science is needed to advance the field while developing the quantum industry ecosystem.

Managing this delicate equilibrium requires robust EU and member states fundamental research funding, requiring that resulting publications and datasets be openly accessible. This ensures a steady flow of new discoveries leveraged for the development of commercial ventures. It goes with an effective intellectual property (IP) framework that allows researchers to publish freely while securing intellectual property and the launch of spin-offs startups. Moreover, adopting an open innovation mindset, in which pre-competitive platforms and standards are developed collaboratively, enables companies to compete on final products, services and business models, while benefiting from shared R&D. Academic–industry partnerships, joint research centers, and private-sector sponsored fellowships can contribute to enriching the talent pool and nurture knowledge transfer without locking valuable insights behind corporate walls. Working on standardization and benchmarking tools can also contribute to shape the competitive landscape.

EU funding instruments can contribute to this synergy with individual (ERC) and collaborative research grants (EU Quantum Flagship), proof-of-concept grants (EIC), dedicated incubators and large-cap funding (EIB).

7.7 - Education

The computer science education landscape is currently dominated by artificial intelligence. In order to enable the future programs related to this white paper, EU member states will need to train more scientists at the crossroads of AI and quantum computing. New curriculum should be proposed to develop the skills of quantum-AI scientists and engineers.

7.8 - Societal challenges

Quantum technologies and artificial intelligence are strategic technology fields for the EU. They are now widely addressed at dedicated industry fairs and investor meetings. While quantum technologies themselves can be considered as neutral tools, the societal challenges largely arise from how these technologies are applied. The EU is already tackling the societal impact of these emerging technologies. Indeed, the General Data Protection Regulation (GDPR) and the AI Act regulate many aspects related to data protection, transparency, and accountability in high-tech fields, including QT applications.

Furthermore, existing initiatives in Europe are addressing the societal challenges of QT such as the Quantum Delta Centre for quantum and society in the Netherlands, the QuantWorld project in Germany, the Innsbruck Quantum Ethics Lab in Austria or the Humanities for quantum sciences lab in France.

As joint efforts work best when parties coming from different disciplines can develop a common language, the expertise of societal sciences and humanities (SSH) may play a significant role in the joint development of artificial intelligence and quantum technologies. It will ensure that societal and ethical considerations are embedded within the broader strategic development of quantum technologies in Europe.

8 - List of contributors

Andris Ambainis (University of Latvia)

Antonio Macaluso (German Research Center for Artificial Intelligence)

Daniele Binosi (Quantum Community Network)

Frank Wilhelm-Mauch (Saarland University & Forschungszentrum Jülich)

Giovanni Acampora (University of Naples Federico II)

Gorka Munoz Gil (University of Innsbruck)

Iordanis Kerenidis (CNRS)

Jens Eisert (Free University of Berlin)

Kristel Michielsen (Forschungszentrum Jülich)

Leonardo Banchi (University of Florence)

Olivier Ezratty (EPITA, Quantum Energy Initiative)

Pallavi Bhardwaj (SAP SE)

Philipp Slusallek (Saarland University & DFKI)

Philippe Grangier (CNRS)

Richard Küng (Johannes Kepler University Linz)

Roman Orus (Donostia International Physics Center)

Sabrina Maniscalco (University of Turku)

Tommaso Calarco (Forschungszentrum Jülich & University of Bologna)

Vedran Dunjko (Leiden Institute of Advanced Computer Science)

Acknowledgements

Artur Garcia (Barcelona Supercomputing Center)

Laure Le Bars (SAP SE)

Thierry Debuisschert (Thales Group)

Coordinating team

Enrique Sánchez-Bautista (Quantum Community Network - Brussels Office)

Monica Constantin (Quantum Community Network - Brussels Office)

Ziyad Amodjee (CNRS)

The contributors also acknowledge the authors of the Machine learning for quantum manifesto, as part of the text in the AI for quantum section were drawn from there:

https://www.openpetition.eu/petition/online/support-the-machine-learning-in-quantum-science-manifesto-2#petition-main.

9 - List of acronyms

AI: Artificial intelligence

CFN: Collision-free navigation

EIB: European Innovation Bank

EIC: European Innovation Council

ERC: European Research Council

FTQC: Fault-tolerant quantum computer

HHL: Harrow-Hassidim-Llloyd, a quantum algorithm solving linear problems

HEP: High-energy physics

HPC: High performance computing

LLM: Large language model

MAS: Multi-agent system

MDP: Markov decision process

ML: Machine learning

NISQ: Noisy intermediate-scale quantum

NN: Neural network

PCA: Principal Component Analysis

POMDP: Partially observable Markov decision process

PQC: Parametrized quantum circuit, also “post-quantum cryptography” but not in the context of this white paper.

QAOA: Quantum approximate optimization algorithm

QCNN: Quantum convolutional neural network

QEC: Quantum Error Correction

QEM: Quantum Error Mitigation

QFT: Quantum Fourier transform

QHD: Quantum Hamiltonian descent

QMAS: Quantum multi-agent system

QMDP: Quantum Markov decision process

QML: Quantum machine learning

QNN: Quantum neural network

QP: Quantum-supported planning

QPS: Quantum planning and scheduling

qRAM: quantum random access memory

QRL: Quantum reinforcement learning

QSL: Quantum supervised learning

QS: Quantum-supported scheduling

QT: Quantum technology

QUBO: Quadratic unconstrained binary optimization

RL: Reinforcement learning

SRIA: Strategic Research and Industry Agenda

VQA: Variational quantum algorithm