Torta et al. *EPJ Quantum Technology* (2025) 12:66
https://doi.org/10.1140/epjqt/s40507-025-00369-8

EPJ.org

EPJ Quantum Technology
a SpringerOpen Journal

REVIEW

Open Access

Quantum computing for space applications: a selective review and perspectives

Pietro Torta[1*], Rebecca Casati[1], Stefano Bruni[1], Antonio Mandarino[1] and Enrico Prati[1*]

*Correspondence:
pietro.torta@unimi.it;
enrico.prati@unimi.it
[1]Department of Physics Aldo Pontremoli, Università degli Studi di Milano, Via Celoria 16, Milan, 20133, Italy

Abstract

Space science and technology are among the most challenging and strategic fields in which quantum computing promises to have a pervasive and long-lasting impact. We provide an overview of selected published works reporting the application of quantum computing to space science and technology. Our systematic analysis identifies three major classes of problems that have been approached with quantum computing. The first category includes optimization tasks, often cast into Quadratic Unconstrained Binary Optimization and solved using quantum annealing, with scheduling problems serving as a notable example. A second class comprises learning tasks, such as image classification in Earth Observation, often tackled with gate-based hybrid quantum-classical computation, namely with Quantum Machine Learning concepts and tools. Finally, integrating quantum computing with other quantum technologies may lead to new disruptive technologies, for instance, the creation of a quantum satellite internet constellation and distributed quantum computing. We organize our exposition by providing a critical analysis of the main challenges and methods at the core of different quantum computing paradigms and algorithms, which are often fundamentally similar across different domains of application in the space sector and beyond.

Keywords: Quantum Computing; Quantum Annealing; Quantum Machine Learning; Space Science and Technology; Scheduling problems; Earth Observation; Quantum Technologies

1 Introduction

Among many fields anticipated to benefit from quantum computing, space science and technology have a special role. High-performance computing and intensive computational resources are essential for a wide range of applications. Some notable examples include space mission scheduling through efficient task planning, processing large datasets of airborne or satellite images for remote sensing, and simulating fundamental physics processes. This review provides an overview of research conducted at the intersection of quantum computing and space science and technology.

Although we aim to discuss how quantum computing [1, 2] entered the quest to solve the computational challenges posed by space activities, our focus is on quantum computing architectures and algorithms, which are discussed concisely but rigorously. We hope that

Springer

this approach may help to elucidate the core challenges and opportunities of near-term quantum computing, which often transcend the specific domain of application. To assist readers – especially those with a technical background in the space sector – in navigating our Selective Review, we have included a summarizing Table in Sect. 5, which organizes all reviewed papers on quantum computing for space science and technology according to the classification framework of the Jet Propulsion Lab of NASA.

With this review, our objective is to provide a timely and thorough assessment of current state-of-the-art and future directions in this interdisciplinary field. Previously, NASA, which has pioneered quantum computing applications in space, has periodically publicly shared the status of its advancement on quantum annealing [3, 4] and gate-based quantum computing [5]. However, we feel that a review that gathers the latest developments is missing.

We present a focused summary of relevant literature, structured into three primary research areas, each covered in a dedicated section. A visual synopsis of the reviewed subjects is provided in Fig. 1, while detailed summaries are included at the beginning of each section. We begin with Sect. 2, which examines QUBO problems in space science and technology, primarily tackled with quantum annealing (QA) or hybrid quantum-classical schemes based on QA. Particular emphasis is placed on scheduling problems, with a discussion of their mathematical formulations and their relevance to space applications.

Next, Sect. 3 explores the use of quantum machine learning (QML) algorithms for remote sensing and Earth Observation. The need for processing and classifying vast datasets from airborne and satellite systems has recently led to numerous quantum and hybrid quantum-classical algorithmic proposals. Most QML algorithms are designed to run on processors supporting a gate-based architecture, by leveraging both quantum and classi-

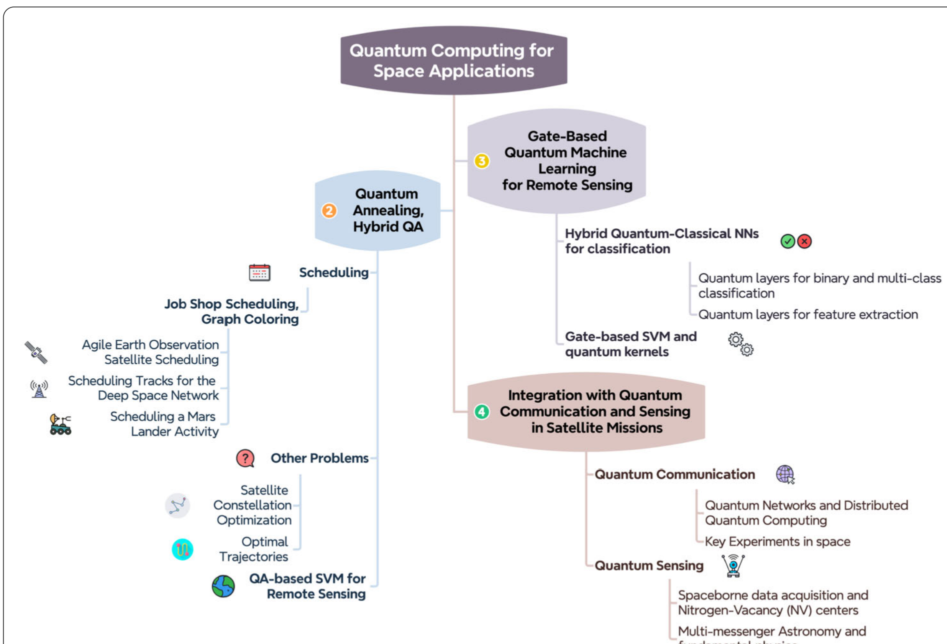

Figure 1 Pictorial description of the main applications of quantum computing to space science and technology covered in this review article. The colored numbers refer to the corresponding sections

cal computation. A notable exception is the QA-based version of SVMs, which is treated in this section for ease of presentation despite relying on a QA scheme (see Fig. 1).

Although the QA-related techniques for optimization and the broad class of quantum machine learning algorithms serve largely independent domains of application, it is worth noting that there are cross-domain implementations. A first remarkable example is represented by the QA-based SVM mentioned above, to be exploited e.g. in classification problems. Secondly, the Quantum Approximate Optimization Algorithm (QAOA) is one of the most renowned gate-based algorithms for optimization and learning, and has a direct connection to QA, as mentioned in Sect. 2.1. Another noteworthy example is represented by training Boltzmann Machines with QA [6], which is not covered here.

Section 4 broadens the perspective of quantum computing for space by considering the potential integration with other quantum technologies. To maintain focus, we do not span the entire breadth of available and proposed schemes that combine quantum technologies and space research. Such efforts have already been addressed in previous topical and comprehensive reviews, devoted to shaping a roadmap for quantum technologies in and for space [7], reviewing tests of fundamental physics implementable in space along with applications of quantum communication and quantum sensing in space [8], and discussing quantum optical implementations [9]. Here, we narrow our scope to the prospects of integrating quantum computing with quantum communication and quantum sensing for satellite missions.

Finally, in Sect. 5, we provide an outlook on quantum computing for space and concluding remarks. A list of acronyms used throughout this work is reported in App. C.

Among the topics that are not covered by our review, we mention physics-inspired and quantum-inspired classical algorithms, based e.g. on tensor networks, which have been analyzed in a recent publication [10]. To provide a coherent selective review that highlights key concepts and challenges in this field, we were unable to incorporate all available published studies with the same level of detail while maintaining a reasonable manuscript length. Consequently, certain references of major impact are examined in greater detail than others, and our review is organized and structured around classes of quantum and hybrid quantum-classical algorithms.

It is our goal to offer space scientists and technologists a practical guide to understand the basics of quantum computing. In parallel, our goal is to establish a solid foundation for further incremental and groundbreaking work in the field. Finally, a quantum computing expert with knowledge of quantum theory and algorithms may find several real-world problems to benchmark their new algorithmic proposals, which are often tested only in idealized and synthetic scenarios.

The remainder of this Introduction offers a concise overview of quantum computing, aimed at readers from different backgrounds.

1.1 Quantum computing in a nutshell

Unlike conventional microprocessors, which rely on a single technology – the silicon platform – and one computational architecture based on Boolean logic implementing the Von Neumann approach, at least three computational architectures enable universal quantum computing.

In the circuit model [11], also known as gate-based architecture, registers of initialized qubits – the quantum computational units – are processed by one-qubit and two-qubit quantum logic gates. Each gate is realized by a sequence of physical operations, implementing approximately unitary operators. The sequential application of these gates results in a quantum circuit, i.e., a multi-qubit unitary transformation. A final (probabilistic) measurement of the qubit register provides the quantum algorithm's answer to the computational problem. Usually, a series of quantum register preparation operations and measurements (often dubbed "shots") are needed to gather statistics on some observable quantity to be measured, described by a hermitian operator.

In the adiabatic quantum computing architecture [12], instead, the physical qubits are exploited as a graph of computational units onto which the computational problem can be mapped, resulting in a complex energy potential with several minima. Initializing the system in the ground state of a simple initial Hamiltonian, it is next driven very slowly (i.e. adiabatically) to the target state, encoding the solution to the problem, which is the quantum state associated with the minimum energy for a final target Hamiltonian.

Lastly, in the measurement-based (or one-way) architecture [13], the system simulates a quantum circuit by starting from a highly entangled state, called a graph state, or resource state, or cluster state. Quantum computation is performed by successive one-qubit measurements of individual qubits, whose output is used to condition the next qubits until a final register containing the output of the algorithm is measured to obtain the answer. This makes it an irreversible type of computation from its core principles.

From a theoretical perspective, all three computational architectures are polynomially equivalent and, therefore, quantum computationally universal. However, their experimental implementations are currently at very different stages of maturity. Six major quantum computing prototype technologies exist [14], namely superconductors [15], trapped ions [16], neutral atoms [17], spin in silicon [18–20], color-centers in diamond [21], and light states in photonic circuits [22]. They all share the capability to encode quantum information by controlled quantum states and to leverage quantum resources like superposition and entanglement. The state-of-the-art of these technologies is rapidly evolving, and a brief overview with supplementary references is provided in App. A.

All universal quantum computation implementations currently rely on the gate-based architecture, with trapped ions and superconducting platforms offering the most advanced Noisy Intermediate-Scale Quantum (NISQ) devices [23]. The adiabatic quantum computation paradigm has inspired a generation of quantum annealers [24], which implement a dynamical evolution of an Ising spin glass by tuning a dynamic coupling with a transverse field. This approach is not universal and natively solves only Quadratic Unconstrained Binary Optimization (QUBO) problems. The leading technology is based on superconducting qubits and may offer the first commercial applications [25]. Finally, the measurement-based architecture may be implemented on photonic devices, but it is still in its infancy.

As the existing hardware of quantum computers is, at the time of writing, still relatively immature, current research on quantum algorithms – and interest in their application to space, as in other industrial fields – is conducted under the assumption that, at some point in the future, the available resources and reliability will be sufficient for practical use. Therefore, although there is currently no advantage in encoding and solving problems with quantum hardware compared to classical methods, the substantial effort invested in developing several distinct technologies is progressively narrowing the gap between available resources and the scale of realistic problems. At the same time, the reliability of qubits has improved dramatically – from a few-qubit processors in the 2015-2023 period – to thousands of qubits and partial quantum error correction in some hardware starting in 2024. Like other fields, space applications will progressively benefit from intermediate advances in hardware on the path toward universal fault-tolerant quantum computing.

There is currently no universal consensus on the terminology used to describe the anticipated advantages offered by quantum computing, which have been referred to using various labels, sometimes with imprecise or ambiguous meanings. In the literature, *quantum supremacy* is commonly defined as the point at which a quantum computer can perform a computational task that is practically infeasible for any classical supercomputer, regardless of whether the task has practical relevance. Highly debated claims of quantum supremacy have appeared in the literature, but none have, in fact, addressed computational tasks of practical relevance beyond academic interest. In some policy and media contexts, this term also carries geopolitical connotations, reflecting the so-called "quantum divide" between countries that possess quantum technologies and those that do not.

Quantum advantage, in contrast, typically refers to a quantum computer solving a *useful* problem more efficiently than classical algorithms. A more scientifically precise concept is that of *quantum speed-up*, namely a reduction in computational complexity – often in asymptotic terms – compared to the best-known (or a class of) classical algorithms. A clarifying and scientifically sound definition of different kinds of quantum speed-up was offered long ago in Ref. [26]. Nevertheless, a quantum advantage may also arise in different forms, since it may more broadly encompass any practical benefit in performance, accuracy, or resource use. An example may be represented by improved energy efficiency [27], or co-integration with other quantum technologies such as quantum sensing and quantum communication.

Quantum utility is a more recent term [28], generally used to describe scenarios in which quantum computation emerges as a valid alternative or complement to classical techniques for a practically relevant task, even in the absence of a proven computational advantage. As this concept is still evolving, its definition can vary depending on the context in which it is used. In some cases, it refers to a quantum computer that is able to run a quantum circuit for a *useful* task, more efficiently than any classical simulation of the same circuit. However, this would not preclude the existence of an intrinsically more efficient classical method.

2 Quantum Annealing (QA) for scheduling and other key challenges

This section covers QA and hybrid quantum-classical techniques based on QA, with a focus on their deployment in space science and technology.

The QUBO formulation is introduced in Sect. 2.1, and the underlying principles of QA are discussed in Sect. 2.2. Practical implementation on D-Wave quantum devices is explained in Sect. 2.3, focusing on embedding challenges, limitations of current quantum hardware, and hybrid solutions allowing to tackle larger and more diverse classes of problems. Section 2.4 offers an updated and critical perspective on the state-of-the-art of QA, which may serve as a solid starting point for practitioners, highlighting common trends and challenges.

Next, we focus our discussion and topical review on scheduling problems, emphasizing recent research and recurring themes. Indeed, several tasks of interest in the aerospace sector can be mapped to scheduling problems, such as Graph Coloring (GC) [29, 30] or

Job-shop Scheduling Problem (JSP) [31]. As briefly discussed in Sect. 2.5, these are two prominent examples of NP-complete problems that can be recast into the QUBO form. Section 2.6 reviews published work on QA for scheduling problems in the space sector, including the satellite scheduling problem, the Deep Space Network scheduling, and the example of a Mars lander activity problem.

Finally, Sect. 2.7 describes a few selected examples of QUBO problems beyond scheduling, yet relevant for space science and technology, and tackled with QA techniques. A few examples include satellite constellation optimization, and determining optimal trajectories for aerial and space vehicles.

2.1 Quadratic Unconstrained Binary Optimization (QUBO)

Space applications pose formidable computational challenges that often admit a reformulation in terms of combinatorial problems, specifically within the subclasses of optimization and decision problems. In addition to the intrinsic interest in solving real-world use cases for specific applications, these computational tasks often fall into renowned families of computer science problems, frequently belonging to the NP-complete or NP-hard complexity classes. These tasks are computationally intractable, as finding an exact solution in worst-case scenarios requires computational resources that grow at least exponentially with the size, and no efficient exact algorithm exists. However, extensive research on heuristic or meta-heuristic algorithms has been developed to tackle instances that are highly relevant in domain-specific applications. In practice, approximate or even exact solutions are routinely obtained with less than exponential resources but without any performance guarantee. Due to this intrinsic hardness, alternative computational paradigms have emerged as promising candidates for solving hard combinatorial problems and may substantially improve the state-of-the-art of standard heuristic algorithms. Quantum computing is a leading candidate in this endeavor, encompassing both the digital gate-based models of quantum computation and analog Quantum Annealing (QA) [24]. The interconnection of different computational problems implies that any numerical or theoretical evidence of some form of quantum speed-up [26] for a domain-specific application would likely imply similar benefits in other areas of application.

The most significant category of combinatorial optimization problems addressable using quantum computing is Quadratic Unconstrained Binary Optimization (QUBO). It consists of the minimization of a quadratic cost function in terms of a set of binary variables $x_i = 0, 1$:

$$E = \sum_{i=1}^{N} \sum_{j=1}^{N} q_{ij} x_i x_j = \sum_{i=1}^{N} \sum_{j \neq i} q_{ij} x_i x_j + \sum_{i=1}^{N} q_{ii} x_i \,, \tag{1}$$

as for such binary variables, it holds $x_i^2 = x_i$. A key observation consists of noticing the formal equivalence of any QUBO problem to an Ising spin glass, one of the best-studied models in statistical physics. An Ising spin glass is defined by its energy function (or Hamiltonian in physical terminology)

$$H = \sum_{i=1}^{N} \sum_{j=i+1}^{N} J_{ij} \sigma_i \sigma_j + \sum_{i=1}^{N} h_i \sigma_i, \tag{2}$$

now expressed in terms of N classical spins $\sigma_i = \pm 1$. This Hamiltonian can be interpreted as a cost function depending on arrays of spins $\boldsymbol{\sigma} = (\sigma_1, \ldots, \sigma_N)$, where only one array (or a specific set of arrays) achieves the minimum energy, corresponding to the ground state(s). The one-to-one mapping between QUBO variables x_i and Ising variables σ_i is obtained by the affine transformation $\sigma_i = 2x_i - 1$.

Several practical challenges admit a conceptual reformulation in terms of a QUBO problem [32] or, equivalently, an Ising spin glass [33]. A selected set of examples relevant to space science is discussed in the following sections. Crucially, Eq. (2) represents the broadest category of problems that current Quantum Annealers can encode. The distinctive feature of an Ising spin glass is the presence of 2-variable couplings only, namely the J_{ij}, without higher-order interactions. Some simple examples of spin glass models have been extensively studied in physics, notably the so-called Ising model in d dimensions, where the coupling matrix (J_{ij}) is sparse, and the only non-zero terms correspond to nearest neighbors spins on a chain (for $d = 1$) or on a cubic lattice (for $d \geq 2$). Typically, these non-zero couplings all assume the same value $J_{ij} = J$. However, for optimization problems that can be mapped into Ising spin glass models, the coupling geometry is not restricted to nearest neighbors. The coupling matrix in these cases is often dense, with elements J_{ij} taking various values, reflecting a more complex interaction structure.

We remark that combinatorial optimization problems, particularly when expressed in the QUBO formalism, have also been routinely studied in the context of gate-based quantum computation, specifically with hybrid quantum-classical algorithms. A special role has been played by the Quantum Approximate Optimization Algorithm (QAOA) since 2014 [34], one of the earliest proposals leveraging the iterative optimization of a PQC conducted by a classical optimizer. This formulation, specifically designed for QUBO problems, has later been generalized and refined in many ways [35–40] and today a whole research field, namely Variational Quantum Algorithms (VQAs) [41, 42], relies on similar ideas. These hybrid quantum-classical schemes are discussed in Sect. 3.2, and utilized in the reviewed literature on QML applications to the space sector. Nonetheless, quantum annealers are more specialized to the class of QUBO computational problems and currently allow for the encoding and the potential solution of significantly larger-scale instances.

2.2 A short theoretical overview on QA

Here, we summarize the theoretical underpinning of QA, directing the interested reader to more specialized material [43–47]. As a first step for the quantum solution of a QUBO problem, one should promote the corresponding Ising classical Hamiltonian to a quantum setting. The standard strategy to obtain a quantum Hamiltonian is to map classical spins in Eq. (2) to quantum spin-1/2 Pauli operators $\hat{\sigma}_j^z$. Hence, the initial cost function is mapped to a quantum Hamiltonian that is *diagonal*, by construction, in the computational basis of quantum computation [1], here identified with the common eigenbasis of all Pauli operators $\hat{\sigma}_j^z$:

$$H(\sigma_1, \ldots, \sigma_N) \to \widehat{H}_z(\hat{\sigma}_1^z, \ldots, \hat{\sigma}_N^z)\,. \tag{3}$$

In a highly idealized scenario, the effectiveness of quantum annealing would rely on Adiabatic Quantum Computation (AQC) [12] performed at zero temperature. In a nutshell, a quantum system is initialized in the easy-to-prepare ground state of a transverse-field non-interacting Hamiltonian

$$\widehat{H}_0 = \widehat{H}_x = -\sum_{j=1}^{N} \hat{\sigma}_j^x \,, \tag{4}$$

and then driven by an external coupling to the ground state of a target Hamiltonian $\widehat{H}_z(\hat{\sigma}_1^z, \dots, \hat{\sigma}_N^z)$, which encodes by construction the solution of a classical computational problem as its ground state. Schematically, the standard way to proceed is to construct an interpolating Hamiltonian

$$\widehat{H}(s) = (1-s)\,\widehat{H}_x + s\,\widehat{H}_z \,, \tag{5}$$

with $s = s(t)$ playing the role of the external coupling. Ideally, the goal would be to pursue adiabatic dynamics, following the instantaneous ground state of $\widehat{H}(s)$, by slowly increasing $s(t)$ from $s(0) = 0$ to $s(\tau) = 1$ in a large total annealing time τ. The main challenge is associated with the large value of τ usually required to avoid exciting higher energy states, which depends on the minimum spectral gap Δ encountered during the annealing, as stated by the adiabatic theorem(s) [48]. This is a subtle theoretical subject reviewed in [12]. In simple terms, the minimum annealing time τ required for adiabaticity typically scales as Δ^{-2}. A narrow or vanishing spectral gap, which is the hallmark of a quantum phase transition in the thermodynamic limit (infinite size of a model), poses a serious challenge, as the required values of τ would grow indefinitely.

On top of the finite values of τ, experimental platforms suffer from other non-idealities, particularly finite-temperature effects and limited qubit connectivity. In practice, an arbitrary QUBO (or Ising) problem[1] is defined by a *problem graph* $\mathcal{G} = (V, E)$ where V is a set of vertices corresponding to *virtual* qubits with a bias h_i, and E is a set of weighted edges, each associated with a non-vanishing coupling J_{ij}. However, this is only the abstract (or virtual) formulation of a QUBO problem, which poses a practical challenge to map it into quantum annealers. Indeed, these are characterized by a certain number of *physical* qubits with connectivity limited only to neighboring sites, which may not match the one of the problem graph, hence the need for non-trivial embedding schemes. The technical aspect of embedding a QUBO problem onto a realistic device is discussed below.

Finally, AQC has been proven universal since it is polynomially equivalent to conventional gate-based quantum computing in the circuit model [49]. However, this is only true by exploiting *non-stoquastic* Hamiltonians, *stoquastic* Hamiltonians being defined as those with only non-positive off-diagonal matrix elements in the computational basis. Current quantum annealers have a scalable design that can only implement stoquastic terms, precisely the Ising spin glass Hamiltonian in Eq. (2) and the transverse-field term in Eq. (4). The implementation of non-stoquastic terms such as positive $\sigma_x \otimes \sigma_x$ interactions is particularly challenging with superconducting flux-qubits technology [50].

[1] Since the two problem formulations are equivalent, they are often used interchangeably. In the following, we shall prefer the Ising formulation, as it is the one directly implemented on quantum annealers (see Sect. 2.3).

2.3 Practical QA: embedding challenges, hardware limitations, and hybrid solutions

This section provides an overview of the current state-of-the-art technology in QA hardware. In our discussion, we primarily focus on the latest superconducting devices from D-Wave, currently providing the most advanced platform for quantum annealing. We also acknowledge alternative technologies capable of executing analog operations on qubit registers, such as neutral atom implementations developed, for instance, by Pasqal [51], Atom Computing [52], and QuEra [53].

D-Wave's Advantage system has a quantum processor with 5000+ superconducting qubits, a significant enhancement compared to the 2000 qubits of its predecessor, since more physical qubits clearly allow to tackle more complex and computationally extensive QUBO problems. This system relies on a Pegasus graph topology, characterized by a 15-way qubit connectivity [54]. Each qubit in the graph can directly interact with 15 neighbors, thus defining a more dense and versatile connectivity map compared to the 6-way connectivity of the precedent Chimera graph topology. This promotes efficient utilization of quantum resources, resulting in a more compact embedding of an abstract problem graph $\mathcal{G}$ into the Quantum Processing Unit (QPU), which typically correlates with improved solution quality. The upcoming generations of quantum annealers are moving towards more complex topologies like the so-called Zephyr with 20-way qubit connectivity, also featuring longer coherence time.

The actual annealing schedule that is experimentally available on D-Wave QPUs is closely related to the theoretical formulation in Eq. (5) and can be formalized as:

$$\widehat{H}(s) = A(s)\,\widehat{H}_x + B(s)\,\widehat{H}_z\,, \tag{6}$$

where the two annealing functions A and B are fixed by the QPU architecture and such that $A(s) \gg B(s)$ at $s = 0$ when the annealing begins, and $A(s) \ll B(s)$ at $s = 1$ at the end of the annealing. The standard linear ramp $s = s(t) = t/\tau$ can be replaced by a piece-wise linear annealing pattern. This flexibility allows the implementation of advanced strategies such as quench and pause or reverse annealing, as mentioned in Sect. 2.4. While the two functions A and B define a global annealing schedule for all qubits, tailored control at the single superconductive qubit level can be implemented, as detailed online in D-Wave documentation [55].

2.3.1 Embedding

The embedding procedure involves mapping the abstract QUBO problem graph, denoted by $\mathcal{G} = (V, E)$, onto the D-Wave quantum annealer. Qubit chains are employed whenever the limited connectivity of the physical qubits in the QPU impedes direct one-to-one mapping. In fact, to represent a high-connectivity graph on the physical hardware, it is possible to encode each virtual qubit using *multiple* physical qubits. These physical qubits should be adjacent in the hardware connectivity map and are artificially constrained to act as a single qubit. In practice, a strong ferromagnetic coupling J_{F} is introduced to artificially force all such spins either in the spin up or down state.

Notably, the Pegasus topology allows for directly embedding triangular relationships among qubits, a significant improvement over the older Chimera graph. Increased hardware connectivity simplifies the embedding process by reducing or eliminating the need to

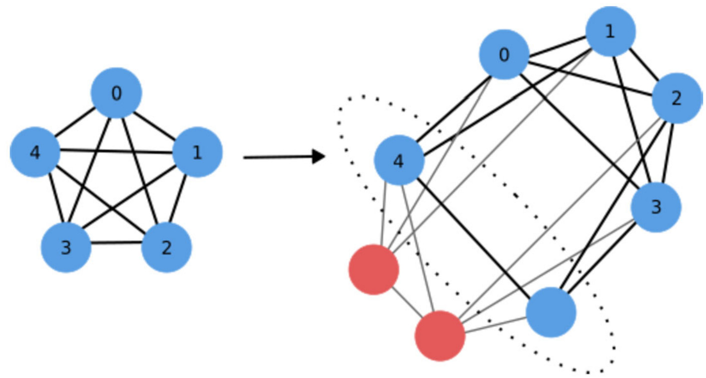

Figure 2 Example of minor embedding of a fully connected graph with 5 virtual qubits, representing a simple QUBO problem, $\min(x_1 + x_2 + x_3 + x_4 + x_5)^2$, onto 6 physical qubits on a subsection of the Pegasus topology. A chain of two physical qubits has been formed to represent the virtual qubit number 4. The red nodes represent physical qubits of the QPU not used for embedding this problem

form chains, potentially enhancing the performance of quantum annealing operations. A comparative study on embedding algorithms in Ref. [56] reveals that the Pegasus topology consistently achieves shorter chain lengths by about 40-60 percent compared to Chimera, significantly improving runtime and embedding efficiency. A simple illustrative example of the embedding process is shown in Fig. 2.

Embedding a problem into a D-Wave quantum annealer is generally a complex task. To facilitate this process, D-Wave provides software libraries that automate the task of finding a suitable embedding for a given problem, with a procedure known as *minor embedding* [57]. On top of reducing the number and the length of qubit chains, a key aspect of the embedding process is setting the "chain strength" value, defined by the ferromagnetic coupling J_{F}. Strong coupling terms are fundamental to ensure that qubits belonging to the same chain align in the same spin state, faithfully representing a single virtual qubit. On the other hand, high ferromagnetic couplings can induce unwanted phenomena such as freeze-out, where part of the quantum system is locked in a particular suboptimal state early in the annealing process. This makes it harder to explore the solution space effectively, thereby reducing the chance of reaching the true ground state. The optimal choice for J_F resides in this delicate trade-off and is usually determined experimentally or by employing D-Wave tools.

Some techniques enable the transformation of higher-order unconstrained binary optimization (HUBO) problems into a QUBO format, as well as the incorporation of specific constraints or problems involving non-binary variables. However, this process often introduces a significant overhead in the number of virtual qubits required, which can complicate the embedding process into the QPU. Various techniques have been developed to address these challenges, as reported in Ref. [25] and references therein.

2.3.2 Hardware limitations

As described in the previous paragraph, only very specific problem graphs are amenable to direct embedding into the QPU. The embedding procedure is computationally hard and may require significant resource overhead, especially for high-connectivity problem graphs.

Even after a correct embedding of a QUBO problem is performed, it is essential to realize that D-Wave is not performing ideal Adiabatic Quantum Computation. Thermal fluctuations impact real machines, as temperatures within D-Wave's freezers typically fluctuate around a few millikelvins. These fluctuations may cause the system to transition to the first excited states during annealing. Furthermore, quantum annealers, similarly to other quantum and analog computing devices, are susceptible to non-idealities such as crosstalk, device variation, and environmental noise.

A shortcoming inherently present in experimental implementations of QA for real-world problems is the lack of availability of prior information about the minimum energy gap Δ encountered along the annealing path. In practice, users have a constrained time interval access to the annealer, typically set between 20 and 2000 microseconds. Obtaining even an order-of-magnitude estimate for the annealing time τ required to maintain adiabaticity is typically infeasible for most practical applications. Short of theoretical guarantees that the available annealing time is sufficient to adhere to the adiabatic theorem, current quantum annealers are rather operated as a *sampler*. At the end of the annealing, each qubit undergoes measurement, and the resulting spin configuration is recorded in a classical binary string. The whole annealing and measurement process is repeated multiple times (usually between 10^3 and 10^4), and a histogram depicting the energy distribution of the measured strings is obtained. Under favorable conditions, the histogram may have a left tail overlapping with the ground state energy of the initial QUBO problem or at least may provide some high-quality approximate solutions.

2.3.3 Hybrid solver service

Despite the steady technological progress in hardware manufacturing, many real-world problems are still too large to fit into current quantum processing units and be solved with QA. This is due both to the intrinsic complexity of practical use cases and to the large number of binary variables – corresponding to virtual qubits – often required to reformulate them as QUBO problems. An example of this mapping to QUBO is summarized in Sect. 2.5 for the graph coloring and the JSP. Moreover, as detailed above, the need for embedding and the use of qubit chains represent an additional computational overhead in terms of physical qubits. In practice, QUBO instances involving more than a few hundred virtual qubits are challenging to embed in the QPU, especially for problem graphs with high connectivity.

A valid alternative strategy relies on hybrid quantum-classical computation. In short, a classical heuristic algorithm decomposes large problem instances into smaller sub-problems that the D-Wave QPU can directly address. These sub-problems are tackled by classical heuristic solvers running in parallel on state-of-the-art CPU and/or GPU platforms. In a nutshell, each node contains a classical heuristic module that explores the solution space and a quantum module, which formulates quantum queries that are sent to a back-end Advantage QPU. Replies from the QPU guide the heuristic module toward more promising areas of the search space or to improve existing solutions. Each heuristic sends tentative solutions to the front-end before a time limit is reached, and the front-end forwards the best results to the user. This approach is hybrid because it uses a classical machine for the initial heuristic decomposition of the problem and the organization of the results provided by the QPUs. Unlike other quantum-classical algorithms, such as Variational Quantum Algorithm, classical computation is not employed as an optimization subroutine.

D-Wave has developed a Python framework [58] that allows user-defined integration of classical and quantum methodologies for the solution of optimization problems, so that the user has full control over the quantum-classical workflow. Alternatively, D-Wave offers a Hybrid Solver Service (HSS), a set of ready-to-use hybrid solvers automatically capable of handling inputs significantly larger than those the QPU can manage alone. This service can be easily employed by users with less technical expertise. Hybrid solvers can also solve a broader set of problems that are not necessarily in the QUBO formulation. To date [59], it is possible to directly solve binary quadratic models (BQM), discrete quadratic models (DQM) and constrained quadratic models (CQMs) without needing to map them into the usual QUBO form. A very recent update includes a novel Nonlinear-Program Hybrid Solver, which is described and tested in Ref. [60]. This article also provides a valuable comparison of the D-Wave hybrid solvers included in the HSS, alongside a selection of previously published works that utilize these hybrid approaches.

2.4 A perspective on QA effectiveness and advanced methods

Many of the most challenging computational problems in industry and science are tackled by heuristic and meta-heuristic algorithms. Despite lacking formal efficiency proofs, remarkable practical performance has been shown by empirical testing and benchmarking on real world use cases. The question of whether QA may provide a genuine speedup [26] over classical methods for QUBO remains largely unanswered, mainly due to the limited problem sizes that current quantum annealers can handle and the inefficiencies in the embedding process.

Hypothetically, QA and hybrid quantum-classical approaches may provide advantages over classical heuristics in different ways. For instance, they could provide higher-quality approximate solutions, faster time to solution and lower energy consumption, or qualitatively different solutions. However, extensive experiments with improved hardware and larger QPUs are mandatory to assess the ultimate effectiveness of QA. This is the conclusion of virtually all papers focusing on QA applications for QUBO with current hardware, specifically those summarized in Sects. 2.6 and 2.7.

Incidentally, it is relevant to note that quantum annealers have also been employed as samplers for different problem classes. A notable example is machine learning tasks, see, for instance, Ref. [4] and references therein, or the more recent work in Ref. [6, 61], including demonstration of *limited quantum speed-up* [62] in an anomaly detection task over real IP traffic data [63]. Potential evidence for quantum advantage has recently been demonstrated for sampling quantum many-body states after a quench (sudden short-time annealing evolution) [64]. While an overview of these topics is beyond the scope of our current work, similar techniques may also be developed for space applications.

To partially overcome the limitations of current QA hardware, recent research has focused on enhancing the algorithmic pipeline, often through the use of hybrid quantum-classical computation. Despite encouraging evidence on the effectiveness of the ready-to-use HSS framework [59], a potential drawback is the difficulty in assessing how much of the empirical success of solving a given task is ultimately due to quantum mechanical effects.

In addition to the HSS, tailored hybrid approaches are frequently designed to address specific tasks. A first strategy resides in carefully preprocessing the QUBO problem, e.g. by pruning unnecessary binary variables and reducing the problem size or utilizing empirical techniques to focus on more relevant sub-portions of the search space. Different post-processing methods on the samples obtained from D-Wave are often used, sometimes in combination with machine learning techniques. Secondly, D-Wave technology allows for some flexibility in the schedule optimization, namely the customization of the schedule $s = s(t)$, which has to be a piece-wise linear function (see Eq. 2.3). The standard linear schedule can be modified by adding a pause during the annealing, which might prove beneficial in some settings [65]. An alternative strategy is reverse annealing [66]. There, the system is initialized in a classical state acting as a warm-start, and the transverse field term is gradually re-introduced and eventually lowered again. This approach might allow for exploring a promising sub-region of the search space with quantum fluctuations, and it can be combined with iterative techniques. Hybrid schemes or schedule optimization have been applied in most publications described in Sects. 2.6 and 2.7.

2.5 Formalization of scheduling problems suitable for QA

Over the past few decades, scheduling problems have received substantial attention in research and industry, with many applications in manufacturing and production environments [67]. A common mathematical formulation is represented by the Job-shop Scheduling Problem (JSP), a fundamental problem in operations research and computer science, which formalizes a variety of scheduling tasks of crucial importance for the space sector. The objective is to schedule jobs on a set of machines (or resources), where each job consists of several operations (tasks) that need to be performed in a specific order. Each operation is characterized by a certain duration and should be assigned to a specific machine, or to a subset of machines [67]. An optimal schedule should minimize the total completion time or maximize a score while meeting operational constraints.

Before diving into the JSP formulation in Sect. 2.5.2, we deemed it instructive to summarize the Graph Coloring (GC) problem [68] in Sect. 2.5.1 for two reasons. First, it may be regarded as a reduced formulation of the JSP, and second, it offers a prototypical example of a combinatorial optimization problem that can be easily recast into the QUBO form.

2.5.1 Graph coloring problem

The concept of graph coloring originates from the practice of coloring adjacent countries on a map, namely a planar graph, with different colors. In modern terms, it is formulated as vertex coloring for a graph $\mathcal{G} = (V, E)$ where V is a set of vertices and E is a set of edges, each connecting a pair of vertices. The K-coloring problem requires finding a proper labeling of the graph's N vertices utilizing K colors. The requirement is that each pair of vertices (i, j) connected by an edge should have different colors .[2] For $K \geq 3$ colors, it is an NP-complete problem. The chromatic number of a graph, often indicated as $\chi(\mathcal{G})$, is the smallest number of colors required to achieve a proper vertex coloring. The determination of $\chi(\mathcal{G})$ for a generic graph is itself an NP-hard problem. The four-color theorem offers a remarkable exception for planar graphs, since a proper coloring of any map with at least $K = 4$ colors exists, and it can be found in polynomial time. An example of graph coloring for a non-planar graph is shown in Fig. 3.

[2] The choice of colors is arbitrary, so one can use integer labeling. A subset characterized by the same color is an *independent set* of the graph $\mathcal{G}$. The face coloring of a planar graph can be recast to the vertex coloring of its dual, and similarly, the edge coloring of a graph is just the vertex coloring of its line graph.

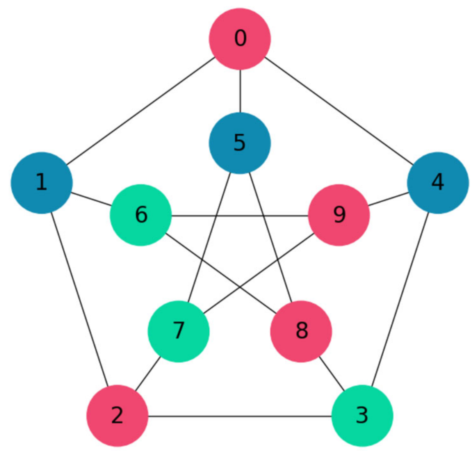

Figure 3 Example of a correctly solved graph coloring problem. This is a Petersen graph, an undirected graph with $N = 10$ vertices and 15 edges. The colors are $K = 3$, which is equal to its chromatic number

Exact methods for the solution of the K-coloring problem require exponential resources in the graph size, and heuristic or meta-heuristic algorithms are employed for realistic scenarios. Among these, quantum and hybrid quantum-classical optimization techniques might prove useful. Indeed, graph coloring can be reformulated as a QUBO problem in terms of a set of binary variables $\{x_{i,c}\}$, where $x_{i,c} = 1$ if the i-th vertex has a label (color) c:

$$H_{GC} = \sum_{(i,j)\in E} \sum_{c=1}^{K} x_{i,c}x_{j,c} + s \sum_{i=1}^{N} \left(1 - \sum_{c=1}^{K} x_{i,c}\right)^2 . \tag{7}$$

Here, s is a parameter that weights the relative importance of two contributions: The first term formalizes the graph coloring task by penalizing identically labeled adjacent vertices. In contrast, the second term is a *soft* local constraint that penalizes configurations in which vertices have zero or multiple color assignments. Note that this formulation requires KN binary variables (or *virtual* qubits, in a quantum setting).

The connection with scheduling can be intuitively understood as follows [30]. Nodes in GC represent tasks to be performed, while edges account for operational constraints, dependencies, or conflicts between tasks. For instance, an edge between two nodes may indicate that two tasks require competing usage of the same resources, typically the same machine or subset of machines, potentially leading to a scheduling conflict. In this context, colors may represent the allocation of a machine or time slot to a task. If two adjacent nodes are colored equally, a conflict in the schedule is counted.

2.5.2 Job-shop scheduling problem

The Job-shop Scheduling Problem (JSP) offers a more detailed mathematical framework for scheduling tasks and is one of the most extensively studied problems in combinatorial optimization. It is a general paradigmatic constraint satisfaction problem (CSP) modeling the optimal allocation of resources required to execute sequences of operations, subject to constraints on location and time.

The JSP formalization involves a set of N jobs $\mathcal{J} = \{J_1, \ldots, J_N\}$, with a job defined as a totally ordered sequence of operations:

$$J_n = \{O_{n1} \to O_{n2} \to \cdots \to O_{nL_n}\}. \tag{8}$$

Here, $n = 1, \ldots, N$ represents the job index, and each job may be composed by a different number of operations L_n. Each operation of any job, e.g. the j-th operation of job n identified by O_{nj}, is characterized by an integer execution time, often dubbed p_{nj}. Moreover, it must be executed on a prescribed machine selected from a set $\mathcal{M} = \{M_1, \ldots, M_M\}$ with cardinality M.

The objective of the JSP is to schedule all operations on the corresponding machines by selecting appropriate integer starting times for each operation. A valid schedule should satisfy two types of constraints. Firstly, precedence constraints must be respected for each pair of consecutive operations within the same job, as specified in Eq. (8). In other words, each operation must be completed before the next one can begin. Secondly, each machine is limited to processing only one operation at any given time. The goal is to schedule all operations to minimize the total completion time for the jobs, which is often indicated with T and called the makespan.

Benchmark datasets and realistic applications often assume the number of operations per job to be constant and equal to the number of machines, i.e. $L_n = M$. In this case, each job requires each machine only for a single operation. For simplicity, we also limit our focus to this setting, as the generalization is straightforward. Valid schedules are in one-to-one correspondence with the selection of a specific ordering (or permutation) of all N operations (one for each job) to be processed by a given machine. Indeed, once these M orderings (one for each machine) are established, each operation can be assigned the earliest possible start time that respects the precedence constraints within the same job. However, a brute force enumeration would require an unfeasible super-exponential complexity

$$(N!)^M \approx \left(\sqrt{2\pi N}\left(\frac{N}{e}\right)^N\right)^M, \tag{9}$$

where we used the Stirling approximation, and typically M is the same order of N. This leads to astronomical runtime even for small instances, such as $N = M = 10$.

Only recently, an exact algorithm with a better scaling than brute force has been found [69]. This relies on dynamical programming, and despite being exponentially faster than brute force, it still requires super-exponential resources:

$$(p_{max})^{2N}(M+1)^N, \tag{10}$$

where p_{max} is the maximum execution time among all operations. Indeed, the JSP has earned a reputation for being especially intractable, and the best general-purpose classical solvers struggle with instances of moderate size, such as the aforementioned $N = M = 10$. To solve larger instances approximately, heuristic and meta-heuristic algorithms are required. A QUBO formulation of the JSP problem has been proposed in Ref. [31], showcasing preliminary results with QA. Further numerical experiments on D-Wave hardware were performed in Ref. [70], and hybrid quantum-classical schemes were devised in Ref. [71]. The same QUBO formulation can be exploited for simulating the adiabatic evolution on a classical variational manifold [72].

To simplify the notation, let us introduce a collective index $i = 1, \ldots, NM$ to enumerate all operations across all jobs, for example, by first indexing all operations of the first job, followed by those of the second job, and so on. In close analogy to the graph coloring case, we introduce binary variables defined as

$$x_{i,t} = \begin{cases} 1 & \text{if operation } O_i \text{ starts at time t,} \\ 0 & \text{otherwise} \end{cases} \tag{11}$$

We remark that in the JSP framework, time is discretized as $t = 0, \ldots, T$.

To recast the problem into QUBO, we frame it as a decision problem, much like the K-coloring problem. By predefining the timespan value T, we aim to determine whether a valid schedule of all operations within this time limit exists. In order to find the minimal value of T that admits proper scheduling, we may solve the decision problem by iteratively lowering T or introducing a penalty term that favors valid schedules with low T [31].

The QUBO formulation of the JSP decision problem has a Hamiltonian in the following form:

$$H_T(\mathbf{x}) = \eta\, h_1(\mathbf{x}) + \alpha\, h_2(\mathbf{x}) + \beta\, h_3(\mathbf{x})\,, \tag{12}$$

where $\eta, \alpha, \beta > 0$ are user-defined constants that weight three different contributions, and $\mathbf{x}$ is a collective vector for all the binary variables introduced in Eq. (11).

The explicit expression of the three contributions and technical details on the implementation and variable pruning are found in Ref. [31]. In a nutshell, the first term $h_1(\mathbf{x})$ enforces the order of operations in a job, by ensuring that each operation can only start after the previous one within the same job has been completed. The second contribution $h_2(\mathbf{x})$ models the machine constraints, namely the processing of at most a single operation on any machine at any time. The latter is composed of two terms: the first penalizes any schedule where an operation begins on a machine before the previous one is completed, and the second penalizes the simultaneous start of two operations on the same machine. Finally, the third term $h_3(\mathbf{x})$ is a soft constraint introduced in the QUBO formulation to ensure that each operation has a unique start time. It penalizes cases where an operation is assigned multiple start times or none at all. Valid schedules do not violate any of these constraints and therefore have energy zero, while violation of any constraint is signaled by a positive energy.

The JSP instances can also be represented as disjunctive graphs, as described e.g. in Ref [73], in analogy to the graph coloring problem. However, these two graph representations are quite different, as the JSP graph contains both directed (conjunctive) edges between consecutive operations of the same job and undirected (disjunctive) edges between operations assigned to the same machine. A formal solution is obtained by turning each undirected edge into a directed one, namely by ordering the operations assigned to each machine.

Although the JSP formulation is quite general, it can be extended in several ways. For instance, the flexible JSP allows specific operations to be executed on a subset of machines rather than a single one [67]. Moreover, it is sometimes necessary to select only a specific time window for the execution of certain operations or jobs [29].

2.6 Scheduling problems in space science and technology tackled with QA methods

In this section, we present a selection of space applications where the primary computational challenge is a scheduling problem. This is modeled as JSP or a closely related formulation and tackled with QA or hybrid quantum-classical schemes leveraging the D-Wave QPU as a sampler.

2.6.1 Agile Earth observation satellite scheduling problem

Agile Earth Observation Satellites (AEOSs) are operated to acquire images from targeted Earth regions while obeying several maneuvering and operational constraints. They are commonly employed for monitoring disasters, tracking environmental changes, and exploring resources. Usually, satellites are equipped with fixed optical instruments and can rotate on three axes, allowing for flexible and timely data collection.

The AEOS scheduling problem (AEOSSP) involves creating an optimal satellite mission plan based on several competing requests for Earth images, namely acquisition requests, each with a specific commercial value. A total score should be maximized while obeying maneuvering constraints, accounting for the minimum time interval required to rotate optical sensors between two successive image acquisitions scheduled on the same satellite. Additional operational constraints may be related e.g. to ground control, on-board memory or satellite constellations.

This optimization task has been studied extensively for decades and admits both discrete-time and continuous-time mathematical formulations [74]. In the discrete framework, the AEOSSP is formulated as an integer linear programming problem, which is generally NP-hard.

In Ref. [75], a simplified version of the AEOSSP has been recast into a QUBO formulation and tackled with quantum annealing methods. Here, the goal is to optimize the selection of imaging attempts for a single satellite in a highly inclined orbit (typically sun-synchronous at 98°), which provides good coverage of the globe and daily passes over the same Earth region at roughly the same local solar time. The optimization aims at the maximum score while satisfying a series of constraints. It can be recast into a single-machine scheduling problem or extended to a satellite constellation. In this study, the authors created a collection of problem instances that are small enough for implementation on a D-Wave quantum annealer, yet incorporate realistic elements of significant industry-scale problems. The results obtained by the D-Wave 2000Q were compared against those from classical solvers, namely exact algorithms for smaller instances and heuristic approaches for larger problems. In more detail, they compared the time-to-solution of a quantum annealer with that of an exact classical solver for problems of small size. Additionally, they assessed the solution quality for larger instances by comparing the quantum annealer with a heuristic classical optimizer.

Despite comparable performance, current technology did not reveal any practical advantage over classical implementations. Nonetheless, quantum annealers could still compete with classical algorithms that benefited from decades of research.

A few recent publications generalized this quantum optimization scheme, including more satellites, additional degrees of freedom and constraints, and realistic instances from actual satellite missions, see e.g. [76] and references therein. Here, the authors extend the previous numerical simulations by comparing different D-Wave architectures with the

HSS and QAOA simulations. Alternative strategies based on hybrid quantum-enhanced reinforcement learning schemes were explored in Ref. [77]. A proof-of-concept application of reverse quantum annealing for the single satellite problem can be found in Ref. [78]. Finally, the satellite scheduling problem has also been tackled with gate-based hybrid quantum-classical techniques such as the QAOA [79].

2.6.2 Scheduling tracks for the deep space network

The Deep Space Network (DSN) is the National Aeronautics and Space Administration's (NASA) infrastructure managing tracking and communication for deep space missions. The DSN comprises three complexes located in California, Australia, and Spain, each featuring several radio antennas, guaranteeing full sky coverage. The DSN routinely manages tracking, navigation, and data transmission for a few dozen spacecraft missions. The objective is to coordinate multiple antenna access requests for transmitting and receiving spacecraft data and to generate a valid schedule that accommodates all of these requests [80]. A valid schedule should account for viewperiod constraints, namely, the specific time interval when any given spacecraft is visible to an antenna. Additional operational constraints must also be considered, including out-of-service antennas due to maintenance or prioritization of an urgent task over others.

The scheduling of the DSN is a highly complex task that relies on both human interventions and automated solutions. It consists of three successive phases and can take up to a few months. Since it involves a synergy of human planning and computational resources, several efforts have been made to lower the computational cost and increase the level of automation. In practice, the solution to the DSN scheduling problem requires placing as many requests for antenna access – commonly referred to as tracks – as possible within a finite time frame, usually a week, while utilizing a finite amount of computational resources.

In Ref. [81], the authors present a QUBO formulation of this problem, attempting to maximize the number of requests satisfied in the timespan while, at the same time, minimizing the number of conflicts between scheduled tracks. Some tasks may require multiple antennas simultaneously, and long tracks may be split into two shorter segments. This combinatorial problem is an example of the JSP within a *fixed* time window (or timespan) of a week. This paper ambitiously addresses real-world operational data with QA techniques, yet the required number of qubits and connectivity was too large for state-of-the-art quantum annealers in 2022. The authors resorted to D-Wave HSS, which can handle QUBO problems with much larger sizes compared to those directly solved by the QPU. The authors also devised a quantum-inspired custom solver, a fully classical algorithm inspired by quantum techniques. The results obtained with the HSS were compared with a conventional mixed integer linear programming approach: in a shorter run time, the D-Wave HSS was able to satisfy more requests and successfully schedule more hours, resulting to be more effective than the classical counterpart on a specific use case. These results offered a promising perspective on the applicability of hybrid quantum-classical methods for practical scheduling problems. A similar outlook is offered by Ref. [82], where the HSS results are compared against standard classical solvers and deep reinforcement learning techniques.

2.6.3 Scheduling a Mars lander activity problem

An operative Mars lander must perform multiple autonomous activities. Examples include obtaining panoramic pictures of the landscape, measuring Martian weather, sampling the Martian soil, and sending the stored data via communication satellites. Some activities, such as collecting a soil sample, can be divided into subtasks that include precedence constraints. These operations may be conducted only during specific time windows, such as in the presence of sunlight or an unobstructed view of the communication satellites. Moreover, solar panels recharge a battery during the daytime, and a constraint on sufficient battery charge should be satisfied before scheduling any activity.

In the scheme of Ref. [83], the Mars lander can perform one task at a time, each with a specific duration, a time window, a precedence constraint, and a battery consumption rate. As the number of possible activities and constraints grows, finding the best activity scheduling is a complex optimization problem. The authors implemented a hybrid quantum-classical framework, namely a Quantum Annealing Guided Tree search, detailed in Ref. [84]. A quantum annealer is operated to generate discrete *candidate* schedules to a subproblem without battery constraints. These candidate solutions are in the form of binary strings. On the other hand, a classical computer manages the global search tree and checks battery-profile feasibility (a continuous constraint) for schedules produced by the quantum annealer. In practice, the classical algorithm prunes unfeasible branches and focuses the tree search on specific regions of the binary tree.

From a computational standpoint, the Mars lander scheduling problem without the battery constraint is a single-machine scheduling problem, which can be encoded in QUBO form and embedded in the D-Wave hardware. This hybrid quantum-classical approach may look particularly promising since the quantum annealer is used as a *sampler*, and the sampled strings (even those violating constraints, with larger QUBO cost function values) are utilized in a structured fashion, in synergy with a classical algorithm.

This hybrid scheme is benchmarked on a set of simplified Mars lander instances, showing that samples generated by the quantum annealer are effective in guiding the classical tree search. From a practical perspective, no confrontation was performed with a state-of-the-art classical solver since realistic scheduling problems involve larger problem sizes and additional technical issues not included in this analysis.

2.7 Selected QUBO problems in space and aerial sectors beyond scheduling

This section aims to broaden our perspective by pointing out some QUBO problems beyond scheduling that are relevant to the space or aerial sectors and for which QA or hybrid QA techniques have been successfully utilized.

A comprehensive review of the literature goes beyond our scope, which is to provide an insightful overview of a partial yet selected set of publications. Before that, we point the interested reader to specific applications that are not described in our review. A first example is QA for fault detection and diagnosis [3, 4, 85], which is not discussed here. We also cite the application of D-Wave to extract features from a dataset of low-resolution satellite images of airplanes in Ref. [86]. Finally, we acknowledge the application of QA [87] and QAOA [88] to multiple-input multiple-output antenna arrays, a leading technology for enhancing wireless communication systems, including satellite constellations [88].

D-Wave quantum annealers have also been employed to train specific types of machine learning models, namely Support Vector Machines (SVMs), for various applications in the space sector, particularly in remote sensing. Although reformulated as a QUBO problem and solved using D-Wave QPUs, this topic is deferred to Sect. 3, which is dedicated to Quantum Machine Learning (QML) applications in remote sensing. This approach enables us to introduce remote sensing and Earth Observation while also outlining the analogies and differences between QA-based and gate-based implementations of Quantum SVMs.

Some computational challenges in the field of Earth Observation and remotely sensed imagery can be directly mapped to the standard QUBO form and tackled with QA methods. An interesting example is provided in Ref. [89], which proposes a QUBO formulation for two relevant problems: the selection of the most informative bands in hyperspectral images and their binary and multi-class classification for a standard remote sensing dataset. We do not review technical details, although similar problems in Earth Observation are discussed in depth in Sect. 3 and tackled with different approaches.

In the following, we delve into two use cases for QA that we find particularly compelling: satellite constellation optimization and the evaluation of optimal trajectories for aerial and space vehicles.

2.7.1 Satellite constellation optimization

Maximizing the coverage of a targeted Earth region by splitting N satellites into k independent sub-constellations is a real-world optimization challenge that can be formalized as an NP-hard weighted k-clique problem [90]. A weighted graph is defined where vertices represent sub-constellations, vertex weights represent coverage, and edges connect compatible sub-constellations, namely those that do not share any common satellite. In this formalism, any k-clique identifies a unique partitioning of the N satellites into k compatible sub-constellations. The computational goal is to determine the k-clique maximizing average coverage. Realistic problem sizes, featuring a few tens of satellites, become intractable for exact solvers, necessitating heuristic approaches.

This weighted k-clique problem can be cast into a QUBO form – similarly to the graph coloring and the JSP discussed in Sect. 2.5 – including constraints for k-cliques and high coverage. However, the QUBO formulation here is only approximate, and QA runs would frequently yield unfeasible solutions violating some constraints. Moreover, the abstract QUBO graph is fully connected, which can be challenging for the current hardware with limited connectivity. To address these limitations, the authors in Ref. [90] employed a genetic algorithm to reduce the QUBO size, allowing the problem to fit within the quantum annealer. Additionally, they introduced a classical post-processing step to efficiently eliminate violated constraints in a candidate QA solution without decreasing its average coverage. This approach showed promising results in experiments using either a simulator or the real D-Wave hardware. This heterogeneous quantum computing approach, which integrates hybrid quantum and classical computation, could be further refined and expanded, demonstrating its utility beyond the scope of Ref. [90].

Satellite constellation optimization has also been studied in a different framework in Ref. [91]. Here, the computational objective is to minimize overlapping monitoring areas targeted by a constellation of observation satellites. By defining binary variables modeling a conflict graph, the problem is formalized as a maximum weight independent set (MWIS) task, another renowned class of NP-hard problems. The authors define an Ising Hamiltonian (or, equivalently, a QUBO cost function) to model the resulting constraints as penalty terms and propose to utilize QAOA to solve the resulting optimization task.

Finally, it is worth mentioning the recent publication in Ref. [92], focusing on Beam Placement and Frequency Assignment tasks, modeled as clique cover (CC) and graph coloring (GC) problems, respectively. These NP-hard computational problems can be reformulated as QUBO, as shown in Sect. 2.5 for the GC problem. Incidentally, the CC problem is computationally equivalent to the GC problem. To tackle large problem sizes, the authors propose a tailored hybrid quantum-classical approach relying on a Hamiltonian reduction method, claiming performance that it is competitive with, or even superior to, that of commercial solvers. Similar techniques can be extrapolated and refined to tackle other QUBO problems with current D-Wave QPUs.

2.7.2 Optimal trajectories for aerial and space vehicles

Trajectory optimization is crucial in space exploration and satellite operations, where identifying the most efficient trajectories is critical to ensuring mission success, minimizing fuel consumption, and optimizing available resources.

A first notable approach to trajectory optimization is described in Ref. [93] and consists in de-conflicting wind-optimal trajectories for air traffic management. Conflict detection and resolution is a crucial task in the aerial and space sectors. Here, the emphasis is on ensuring the safety of the shared airspace while preventing potential conflicts between scheduled flights. By definition, a conflict is a violation of safety norms and protocols established in the airspace, for instance, when two aircraft get alarmingly close to each other. At the same time, fuel costs and the risk of traffic congestion should be accounted for. Due to the increasing levels of air traffic in the last decades, novel approaches are needed to manage the growing complexity.

A promising algorithm involves starting with wind-optimal trajectories, which are the routes each flight would take to minimize fuel costs assuming no other flights are in the air. Such wind-optimal trajectories – depending only on the origin, destination, departure time, and forecast winds – will conflict with each other and thus need to be de-conflicted. In practice, these trajectories must be perturbed, often by delaying departures.

In Ref. [93], this conflict resolution problem was mapped to QUBO and solved using D-Wave for a simplified yet realistic dataset of transatlantic flights. Only a specific case of perturbation, namely time delays, is considered for numerical experiments. A classical preprocessing step is devised to further reduce the need for quantum computational resources.

The first interesting result is the definition and analysis of a problem graph whose nodes correspond to flights and whose edges represent the presence of at least one potential conflict between a pair of flights. Some typical features of real-world graphs are observed and characterized quantitatively, such as a power law for the distribution of the degrees of vertices. Moreover, the connectivity of the problem graph is a good indicator of the hardness of the problem. For instance, a tree would represent a trivial instance, where the conflicts can be removed straightforwardly.

Upon discretization of the configuration space, the original problem is mapped to QUBO and solved using D-Wave machine 2X, based on the Chimera graph. We notice that the mathematical formulation of the QUBO model has a remarkable similarity to the JSP detailed in Sect. 2.5. As expected, upon increasing the problem size, the D-Wave performance decreases significantly, and finding conflict-free solutions for more than 12 flights became too challenging, suggesting the need for better embedding strategies or more advanced annealing schedules. Although not strictly related to space and still preliminary, this work introduces the concept of de-conflicting trajectories by discretizing the configuration space and analyzing a conflict graph, which could be expanded and refined for use in other applications. Applying new-generation quantum annealers or hybrid quantum-classical methods could improve the results and extend the scope of this approach.

The more recent publication in Ref. [94] offers a different perspective on trajectory optimization with QA. The optimal control problem of a single trajectory is converted into binary optimization in the QUBO form, with the flexibility of including several types of constraints. The trajectory is discretized and approximated with polynomials of a different order, whose expansion coefficients should be optimized. This approach has been formulated for linear (or linearized) dynamic systems but could be extended to more general settings. This methodology has been applied to the lunar landing and rendezvous problems, two relevant examples of dynamical systems that can be approximated with linear dynamics under ad-hoc simplifying assumptions. QA showed promising performance, particularly in relation to computational efficiency and in obtaining solution trajectories that satisfy problem constraints.

3 Quantum machine learning for Remote Sensing (RS) and Earth Observation (EO)

This section reviews Quantum Machine Learning (QML) methods and concepts applied to space science and technology. To simplify the discussion, we narrowed our focus to Remote Sensing (RS) and Earth Observation (EO). These constitute a strategic sector for cutting-edge research and have witnessed several proposals for the application of quantum and hybrid quantum-classical computation. We organize our exposition by focusing on classes of quantum algorithms, which primarily rely on gate-based quantum computation. A notable exception, already mentioned in Sect. 1, is represented by QA-based SVMs.

Before describing a selection of the most relevant quantum algorithms and reviewing previous research, we provide an introduction to the field of RS and EO in Sect. 3.1.

Although a summary of modern quantum computing and hybrid quantum-classical computation is beyond the scope of our work, Sect. 3.2 outlines some of the key features of QML and Parametrized Quantum Circuits (PQCs) that we deem relevant for the next sections. At the same time, we identify key theoretical and algorithmic challenges whose practical solution may significantly improve the state-of-the-art. We direct the interested reader to technical papers to delve further into the theory and application of these topics.

Section 3.3 reviews published literature utilizing Hybrid Quantum-Classical Neural Networks (HQC-NNs) in RS and EO. The exposition is organized by collecting together papers that use quantum layers for a specific task in the workflow (e.g. dimensionality reduction and feature extraction vs. final feature transformation and classification).

We finally move to QA-based and gate-based SVMs in Sect. 3.4. After a concise introduction to the classical SVM, we present the theoretical formulation of both quantum variants and summarize published works within the same application domain.

3.1 Remote sensing and Earth observation

Remote sensing (RS) is the process of collecting data or information about an object, area, or phenomenon from a distance, typically using sensors detecting and measuring electromagnetic radiation (e.g., visible light, infrared, or radar) [96]. The major commercial

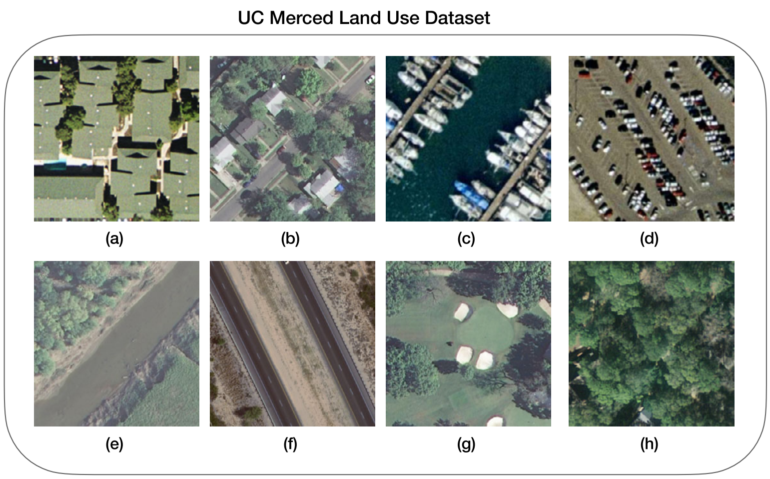

Figure 4 EO images from the UC Merced Land Use Dataset [95], a 21-class land use and land cover dataset containing 100 images per class. Each image is 256×256 pixels with a pixel resolution of 1 foot, with three standard bands Reed, Green, and Blue (RGB). Shown are examples for the following classes, in order: (a) dense residential, (b) medium residential, (c) harbor, (d) parking lot, (e) river, (f) freeway, (g) golf course, (h) forest. Remotely sensed images of certain landscapes can appear quite similar. For instance, a river and a freeway may share similar features, just as a parking lot might resemble a harbor, or a golf course could be mistaken for a forest

and technological application of RS is Earth Observation (EO), which is crucial for mapping, monitoring, and analyzing natural and human-made environments [97]. EO provides valuable data for countless applications, including meteorology and climate change monitoring, disaster prevention and management, agricultural and environmental monitoring, vegetation mapping, and urban planning, among many others [97]. A fundamental task in EO is the identification and classification of land use and land cover images gathered from airborne or space platforms. An illustrative example from the UC Merced Land Use Dataset [95] is shown in Fig. 4.

In the last few decades, technological improvements and increased spatial resolution resulted in a steadily increasing daily volume of satellite images, highly relevant for scientific and commercial missions. In particular, EO satellites are equipped with advanced sensors, including multispectral, hyperspectral, and radar systems, transmitting vast quantities of data, often exceeding dozens of terabytes produced daily. This naturally led to the development of sophisticated algorithms and computational tools to analyze and classify these data, in particular, the widespread adoption of state-of-the-art big data and machine learning tools. An impressive amount of papers have been published on deep learning applications for remotely sensed data [98], and comprehensive reviews and assessments have been conducted, highlighting the progress and challenges within this interdisciplinary research field [99, 100]. A challenge in supervised learning for EO is the mismatch between EO imagery and standard natural image datasets like ImageNet. In fact, remote sensing datasets are often peculiar, and multispectral or hyperspectral images commonly include more bands than the standard Reed, Green, and Blue (RGB). Standard pre-training via transfer learning performed, e.g., on ImageNet, may thus be less effective. The difficulties intrinsically present in EO datasets include low type separability (e.g. rivers and highways can look very similar from a satellite), large differences in the scene scale (the remote sensors may be at different altitudes and have different resolutions), and the coexistence of multiple objects in scenes with different labels (e.g. industrial and residential areas can include similar buildings). In addition to standard RGB, multispectral, and hyperspectral imagery, another powerful remote sensing technology is Synthetic Aperture Radar (SAR). This uses microwave radar to capture high-resolution images of the Earth's surface and can also operate under unfavorable weather conditions (such as cloud cover) and at night. Interferometric SAR (InSAR) extends this capability by analyzing phase differences between multiple SAR images to measure surface deformations and topography with millimeter-scale precision.

As effectively depicted in Ref. [101], EO imagery is characterized by the so-called 4V features: volume, variety, velocity, and veracity. *Volume* simply refers to the massive scale of EO datasets, *Variety* reflects the diversity in spectral information and resolution discussed above, *Velocity* instead is due to the rapid changes observed on Earth's surface over time, and finally, *Veracity* highlights the challenges of dealing with imperfect data, such as noisy images or remote sensing data partially obscured by cloud cover. Perhaps unsurprisingly, the significant computational demands in the fields of RS and EO led to the exploration of alternative computational paradigms, primarily quantum computing. In fact, this is likely the most prolific sector within space science and technology, with the highest number of publications proposing and utilizing quantum or hybrid quantum-classical approaches. We point out the complementary perspective offered in Ref. [102], focusing on computational complexity and the integration of classical high-performance computing and quantum resources, offering at the same time a partial list of publications in the field of quantum computing for RS and EO. From a different angle, QML methods have also been envisaged for addressing computational tasks relevant to the analysis and mitigation challenges related to climate change, as reviewed in Ref. [103]. These challenges are often linked to climate monitoring and modeling via remotely sensed data.

3.2 Parametrized quantum circuits and Variational Quantum Algorithms (VQAs)

The purpose of this section is to provide a concise yet rigorous overview of standard techniques in hybrid gate-based quantum computation, which are employed in the reviewed papers discussed in the subsequent sections.

We do not cover the basics of quantum computing, such as quantum gates, entanglement, and the milestone algorithms in quantum computing theory [1], but rather review hybrid gate-based quantum algorithms that are currently studied and benchmarked on NISQ devices. This is done in Sect. 3.2.1, whereas Sect. 3.2.2 summarizes the critical challenges to resolve in the quest for quantum utility. We aim to offer a practical roadmap for future research on QML for space science and technology, moving from idealized proof-of-principle tests to more realistic benchmarks.

3.2.1 VQA framework and the hybrid quantum-classical optimization loop

Our discussion here is adapted from [104], with additional insights concerning the state-of-the-art of QML for space applications.

The vast majority of hybrid quantum-classical algorithms relying on gate-based quantum computation are essentially devised in terms of Parametrized Quantum Circuits (PQCs) [105], which depend on a set of tunable parameters $\boldsymbol{\theta}$ that can be adjusted to optimally perform a computational task.

These parameters are typically learned using a hybrid quantum-classical approach, in the framework of Variational Quantum Algorithms (VQAs) [41], which are the main candidates for near-term practical applications of gate-based NISQ devices. In fact, VQAs have been devised to jointly leverage the strength of quantum and classical computational resources, to overcome the limitations imposed by NISQ hardware [23], such as noise, limited number of qubits and sparse connectivity, and short coherence time. In essence, the *classically hard* part of the algorithm — i.e. preparing and measuring non-trivial quantum states — is delegated to a gate-based quantum device, whereas the remaining computations, in particular the optimization step that iteratively refines the parameters $\boldsymbol{\theta}$, are performed on a standard classical computer. The optimization step can be done either with gradient-based or gradient-free methods.

The whole framework relies on a quantum-classical optimization loop, synthesized in Fig. 5. Here, we present a simplified formulation that preserves the essential characteristics of a general VQA, applicable to tasks such as preparing the ground state of a quantum system and solving classical optimization problems. It is fundamentally equivalent to the well-known variational principle used in techniques like Variational Monte Carlo. The cost function reads

$$\mathrm{C}(\boldsymbol{\theta}) = \langle \psi_0 | U^\dagger(\boldsymbol{\theta})\, \mathrm{O}\, U(\boldsymbol{\theta}) | \psi_0 \rangle \,, \tag{13}$$

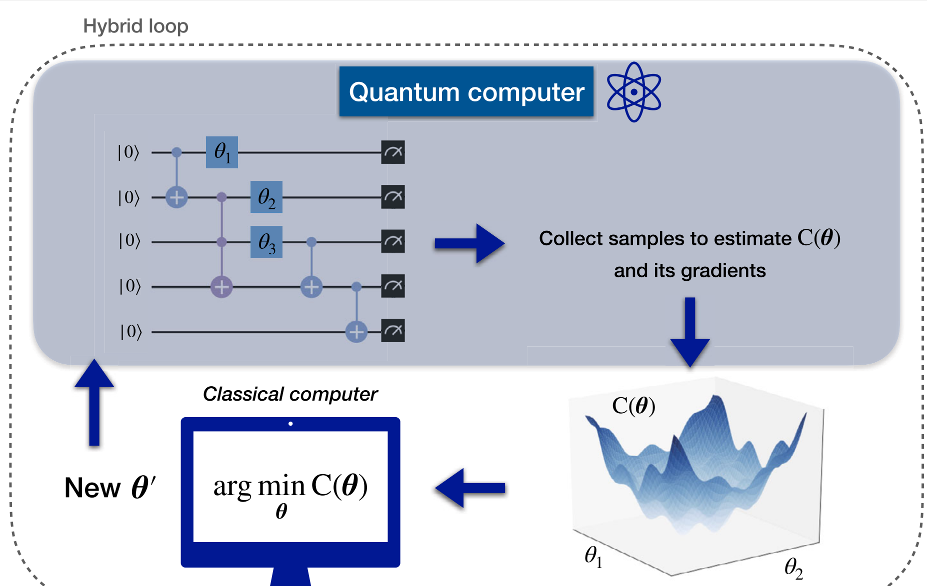

Figure 5 Hybrid quantum-classical loop of a VQA, in the simplified version of Eq. (13) (with an initial blank register). The quantum device is utilized only for classically hard tasks, as depicted in the shaded region. At each iterative step, the variational angles $\boldsymbol{\theta}$ — encoded in the PQC — are fixed: The quantum device is used for *repeated* state preparations and measurements. Once the state is prepared, qubits are measured in some bases, and the outcomes are collected to estimate the cost function $\mathrm{C}(\boldsymbol{\theta})$ and its gradients. A pictorial two-dimensional cut of the cost function $\mathrm{C}(\boldsymbol{\theta})$ is shown: the cost function is non-convex, implying a challenging optimization. The estimated value of $\mathrm{C}(\boldsymbol{\theta})$ (and its gradients) are passed to a classical computer, where an optimization routine performs the step $\boldsymbol{\theta} \to \boldsymbol{\theta}'$, proposing a new, improved set of variational parameters. A new PQC is then repeatedly prepared and measured in the next iteration. Figure adapted from Ref. [104]

where the state $|\psi_0\rangle$ is an easily prepared initial state, and $U(\boldsymbol{\theta})$ is a parametrized unitary. The observable O may coincide with the quantum Hamiltonian whose ground state has to be prepared, or with the classical optimization cost function embedded in a quantum setup (see Eq. (3)).

In practice, the initial state is often a blank register $|\psi_0\rangle = |0\rangle^{\otimes N}$, whereas the parametrized unitary operator $U(\boldsymbol{\theta})$ is constituted by layers of single-qubit and entangling two-qubit gates, some of which are parametric as shown in Fig. 5. These are typically rotational gates, such as $R_j^\alpha(\theta) = \exp^{-i\hat{\sigma}_j^\alpha \theta}$, where $\hat{\sigma}_j^\alpha$ is a Pauli matrix with j being the qubit index and $\alpha = \{x, y, z\}$. Other more complex parametric gates are often used, e.g. in the Quantum Approximate Optimization Algorithm (QAOA) [34]. Nonlocal gates, usually acting on pairs of qubits, are crucial to generating quantum correlation among the qubits and to implement all controlled operations. The most common entangling gate is the CNOT gate, which forms a universal set with local operations: any quantum circuit can be decomposed in a finite sequence of such gates. However, different choices can be made depending on the experimental implementation of the processing unit. As a pivotal example, we mention the Mølmer-Sorresen gate, the first one proposed to generate entanglement in trapped-ions simulators [106].

The crucial feature of a VQA is the synergy between quantum and classical computational resources. At each iteration, the quantum hardware is used to *repeatedly* prepare the PQC with fixed values of $\boldsymbol{\theta}$ and to perform measurements on it, yielding estimates for $C(\boldsymbol{\theta})$ and its gradients, if the latter are needed. These values are then passed to a classical computer, which performs an optimization step by updating the set of variational parameters $\boldsymbol{\theta} \to \boldsymbol{\theta}'$. This procedure is iterated up to convergence, or an exit condition is met. This simplified scheme includes the two earliest proposals of VQAs from 2014: the Variational Quantum Eigensolver (VQE) [107] for quantum ground state preparation, and the QAOA [34] for classical combinatorial optimization.

VQAs have also been devised and adapted to a variety of machine learning tasks, such as classification, regression, generative modeling, and numerous other applications [41]. In the context of Quantum Machine Learning (QML), PQCs are sometimes referred to as Quantum Neural Networks (QNNs), stressing their conceptual similarity to classical neural networks, with quantum gates replacing artificial neurons. Nonetheless, it is worth remarking on the notable differences between these two classes of models. Firstly, the computation of gradients through efficient backpropagation is not easily generalized to the QNN framework. Gradient computation proves more computationally demanding and relies on repeated circuit preparation and measurements, as outlined in Sect. 3.2.2. Secondly, modern deep learning architectures are routinely trained with billions of parameters or more, in an overparametrized regime. In contrast, the effective training of QNNs is currently a significant computational bottleneck, mainly due to unfavorable training landscapes with many local minima traps and vanishing gradients, as summarized in Sect. 3.2.2. Additionally, embedding classical data into a quantum circuit (via quantum feature maps) is a non-trivial design choice that significantly impacts performance, a topic covered below in this section. Finally, since QNNs leverage NISQ devices, decoherence and gate errors may pose practical limitations on their performance.

Here, we are interested in QML applications to EO and RS, which are discussed at length in Sect. 3.3 for Hybrid Quantum-Classical Neural Networks (HQC-NNs), and in Sects. 3.4.4 and 3.4.5 for gate-based SVMs using quantum kernels. Therefore, the previous

VQA scheme should be slightly modified to account for a classical training set (such as remotely sensed imagery). In particular, each image – usually after classical pre-processing – should undergo the crucial step of quantum embedding. This consists of encoding[3] classical data into a quantum circuit, which can be done in several distinct ways. A thorough discussion of basic quantum data encoding techniques can be found in Ref. [108], whereas more application-oriented embedding schemes are ubiquitous in the QML literature. A few basic examples are discussed below, and some others will be mentioned in the next sections when reviewing application papers.

Following Ref. [41], the most general cost function, which holds also for QML applications with a classical training set, can be written in abstract terms as

$$\mathrm{C}(\boldsymbol{\theta}) = f\left(\{\rho_k\}, \{\mathrm{O}_k\}, U(\boldsymbol{\theta})\right), \tag{14}$$

where f is some function, $U(\boldsymbol{\theta})$ is a parametrized unitary, $\{\rho_k\}$ is a set of input (or initial) states that can be interpreted as the training set, and $\{\mathrm{O}_k\}$ is a set of observables. In practice, most implementations of a variational scheme rely on a cost function in the form

$$\mathrm{C}(\boldsymbol{\theta}) = \sum_k f_k\left(\boldsymbol{\theta}, \rho_k\right) = \sum_k f_k\left(\mathrm{Tr}[\mathrm{O}_k U(\boldsymbol{\theta}) \rho_k U^\dagger(\boldsymbol{\theta})]\right). \tag{15}$$

More explicitly, each initial state ρ_k is time-evolved by the parametrized unitary and then used to compute an expectation value for each observable O_k: these expectation values are post-processed by applying functions f_k and then summed together to yield the total cost. At each iterative step of the VQA loop for fixed $\boldsymbol{\theta}$, it is thus necessary to *repeatedly* prepare and measure several quantum circuits, one for each input state $\{\rho_k\}$. The previous simpler version in Eq. (13) is simply recovered by dropping the sum over k, and choosing $f_k(x) = x$, $\rho_k = |\psi_0\rangle \langle\psi_0|$.

Notice that the input states $\{\rho_k\}$, here expressed in the general density matrix formalism, are often *pure* states obtained by encoding a classical input from a training set into a quantum circuit. This is shown below, after briefly describing quantum data encoding and citing the most common techniques.

Embedding classical data into a quantum circuit A crucial task for QML applications consists of the embedding (or encoding) of classical data into a quantum circuit, via a *quantum feature map*. The simplest example refers to a single *integer* data point x in the range $\{0, \dots, 2^{\mathrm{m}} - 1\}$ with its binary representation $x = \sum_{j=0}^{\mathrm{m}-1} x_j 2^j$, where $x_j \in \{0, 1\}$ for all j. This value can be embedded in a quantum circuit via *basis encoding*, which is given by:

$$|\psi_{\mathrm{enc}}\rangle = \bigotimes_{j=0}^{\mathrm{m}-1} |x_j\rangle = |x_{\mathrm{m}-1}\rangle \cdots |x_0\rangle, \tag{16}$$

where $|x_j\rangle$ can be either $|0\rangle$ or $|1\rangle$, which represent the computational basis states of a qubit. This is equivalent to the mapping of classical spins into Pauli matrices $\hat{\sigma}_j^z$ as in Eq. (3). The same integer data point can be encoded via Fourier encoding with multiple

[3] We use the terms encoding and embedding interchangeably.

phase gates [108]. Both methods require m qubits and a layer of single-qubit rotations. Note that the binary expansion may as well describe an approximated real number, e.g. in float precision.

A more realistic situation concerns the embedding of a feature vector $\mathbf{x} \in \mathbb{R}^q$. In principle, each component may be decomposed into an approximate binary representation; however, it is more practical to proceed with different embedding methods. The two basic approaches, which have been modified in several ways [108], are *amplitude encoding* and *angle encoding*.

The former requires prior normalization of the feature vector such that $\|\mathbf{x}\|_2 = 1$, and it also applies to a complex-valued feature vector. Without loss of generality, we can assume $q = 2^{\mathrm{m}}$ by filling the remaining x_i with zeros, if necessary. The corresponding *amplitude encoding* equation is given by:

$$|\psi_{\mathrm{enc}}\rangle = \sum_{i=0}^{q-1} x_i |i\rangle_{\mathrm{m}}, \tag{17}$$

defined on the computational basis states of m qubits, given as usual by the binary decomposition of the *integer* i, i.e. $|i\rangle_{\mathrm{m}} = |i_{\mathrm{m}-1}\rangle \cdots |i_0\rangle$. Here, an exponential (in m) quantity of information is stored in m qubits only; however, the gate decomposition required to implement this encoding has an exponential depth. This depth can be reduced by increasing the number of qubits, with *bidirectional encoding* [108].

Angle encoding simply encodes the feature components into the parametrized Pauli rotations of a qubit, such as

$$|\psi_{\mathrm{enc}}\rangle = \bigotimes_{i=0}^{q-1} R_y(2\theta_i)|0\rangle = \bigotimes_{i=0}^{q-1} [\cos\theta_i|0\rangle + \sin\theta_i|1\rangle] \tag{18}$$

where $\theta_i \in [0, \pi/2]$ are the rescaled feature components and, here, the number of qubits is simply $\mathrm{m} = q$. This approach has been generalized in endless ways in QML, essentially by replacing the y Pauli rotations with different parametrized single or multi-qubit gates. Although this basic formulation uses one qubit per feature component, it can be easily modified e.g. by encoding multiple components in one qubit only.

For any embedding scheme, a gate decomposition should be devised to perform the embedding in practice starting from a blank qubit register. This is done by using an appropriate set of elementary gates – some of which are parametrized by the feature vector – resulting in an encoding unitary operator:

$$|\psi_{\mathrm{enc}}(\mathbf{x})\rangle = U_{\mathrm{enc}}(\mathbf{x})|0\rangle_{\mathrm{m}}, \tag{19}$$

where the number of qubits m depends on the embedding scheme. Indeed, it is crucial to understand that different embeddings require a different number of qubits, gates, and circuit depth. Additionally, each method allows different flexibility in manipulating and retrieving the classical information embedded into the quantum circuit.

When considering a training set with several feature vectors $\{\mathbf{x}_k\}$, we can easily draw a connection with Eq. (15) and the input states $\{\rho_k\}$. The whole PQC is essentially divided into two blocks. The first block $U_{\mathrm{enc}}(\mathbf{x}_k)$ performs the embedding of the classical data $\mathbf{x}_k$

into the quantum circuit as in Eq. (19), hence its parametrized gates are fixed and *not trainable*. The second block $U(\boldsymbol{\theta})$ is instead the *trainable* PQC figuring in Eq. (15), whose variational parameters $\boldsymbol{\theta}$ should be adjusted in the training phase summarized in Fig. 5. In the notation of Eq. (15), ρ_k is simply the density matrix associated with the pure quantum state obtained after the embedding, namely in the *ket-bra* notation:

$$\rho_k = |\psi_{\mathrm{enc}}(\mathbf{x}_k)\rangle\langle\psi_{\mathrm{enc}}(\mathbf{x}_k)| = U_{\mathrm{enc}}(\mathbf{x}_k)|0\rangle_{\mathrm{m}}\langle 0|_{\mathrm{m}} U^{\dagger}_{\mathrm{enc}}(\mathbf{x}_k)\,. \tag{20}$$

This is essentially the proposal of Ref. [109], known as Variational Quantum Classifier, where the whole circuit can be formally written into the subsequent application of an encoding block and then a variational block.

For QML applications, data are loaded with the scope of influencing the behavior of a quantum neural network and enhancing its performance. More complex schemes can implement data reuploading and alternate variational layers parametrized by trainable parameters with encoding layers parametrized by the input data [110]. The generalization of the VQA formalism to this case is straightforward.

As a final note, once data is encoded into the register and a variational quantum circuit is defined, the whole circuit must be *transpiled* into a set of native quantum gates suitable for hardware implementation. The performance of a PQC can significantly benefit from *software and hardware co-design*, which accounts for hardware connectivity constraints and minimizes transpilation overhead. However, these aspects are not addressed in the literature on QML for space science and technology that we reviewed. They represent, therefore, a promising direction for future research.

3.2.2 A perspective on challenges in gate-based quantum machine learning

In recent years, the theory of hybrid quantum-classical computation has seen remarkable progress, with numerous algorithmic proposals and steady experimental advancements. However, clear evidence of a practical quantum speedup over the best classical algorithms remains elusive or highly debated across all application domains.

Here, we wish to remark on some theoretical and algorithmic challenges that are very rarely addressed in papers focusing on applications, in particular in the literature on gate-based QML for space science and technology. Nonetheless, these are of crucial importance in the quest for any quantum speed up or quantum utility.

The first one is an almost-trivial remark, which still represents a key direction of improvement. Let us consider, for ease of notation, the simpler case in Eq. (13) (the generalization to Eq. (15) is straightforward). Despite small-scale classical simulations are often run by exactly computing the PQC, by employing dedicated software likewise *Qiskit* [111] or *Pennylane* [112], with real hardware it is not possible to access the exact values of the cost function, which should be estimated as a sample mean (empirical mean estimator) over a certain number of "shots" K:

$$\mathrm{C}^{(\mathrm{K})}(\boldsymbol{\theta}) := \frac{1}{\mathrm{K}}\sum_{j=1}^{\mathrm{K}} \mathrm{C}_{\mathrm{j}}(\boldsymbol{\theta})\,. \tag{21}$$

We remark that the right-hand side of the previous equation represents the average over empirical values of the cost function. Each of these values may require multiple measurements of $U(\boldsymbol{\theta})|\psi_0\rangle$ in different bases, typically determined by the decomposition of the operator O as a weighted sum of Pauli strings. We denote *shot noise* the statistical noise that is inherently present due to a finite value of K, even in the idealized scenario of a perfectly working quantum device without any physical noise or gate errors: by elementary statistics, we expect the statistical error to scale as $K^{-1/2}$. Including at least "shot noise" in the small-scale classical simulation of quantum circuits is a vital step toward realistic applications since exact state-vector simulations represent a simplistic limit of infinite shots. A finite number of shots introduces numerical instabilities in estimates for the gradients, e.g. through variants of the parameter shift rule [113], which are unavoidable in experiments. Clearly, a more realistic simulation would involve incorporating simulated hardware noise or directly implementing the PQC on a real quantum device with physical noise and error rates. This would enable testing error mitigation techniques, which are an essential tool for current quantum computers to successfully solve real-world problems [114].

Our second comment concerns the *trainability* of VQAs, namely the effectiveness of a classical optimization routine in the minimization of the cost function $C(\boldsymbol{\theta})$. The optimization landscape is usually highly non-convex, and *local minima traps* represent a serious issue. A systematic characterization of the presence of local minima traps in the cost function landscape for certain PQCs was derived in Ref. [115], and VQA training was formally proven to be NP-hard in Ref. [116]. Another trainability issue has become a whole new research field, namely *Barren Plateaus* (BPs): in a nutshell, the magnitude of the gradients of the cost function $C(\boldsymbol{\theta})$ tend to vanish exponentially almost everywhere by increasing the number of qubits N. This is equivalent to exponential cost function concentration around a mean value almost everywhere in the optimization landscape [117]. In a nutshell, to accurately estimate gradients $\nabla_{\boldsymbol{\theta}} C(\boldsymbol{\theta})$, one would need an *exponential* number of state preparations and measurements (i.e., shots), at each iteration of the hybrid quantum-classical loop in Fig. 5. This would be impossible beyond a few tens qubits, i.e. precisely in the regime beyond the reach of classical simulations. However, note that small-scale exact state-vector simulations systematically hide the issue of vanishing gradients since they implicitly use an infinite number of shots for a small number of qubits. A short, accessible summary of BP theory, pointing to relevant literature, can be found in [104]. This includes an overview of techniques to mitigate BPs: one of the most promising approaches resides in *warm starts* — i.e. clever initialization strategies for the parameters $\boldsymbol{\theta}$ — which may initialize the search in a region where the gradients are non-vanishing.

To fully avoid the BP problem, the real challenge resides in the definition of a good PQC architecture that must balance expressivity – encoding non-trivial, classically hard quantum states – and trainability, the ability to optimize parameters effectively. Crucially, the PQC should at the same time be *trainable*, and *not classically simulable*, as discussed in [118]. For the interested reader, a unified theory of all possible sources of BPs for generic deep PQCs has been proposed independently by two groups, in Refs. [119] and [120].

We close this perspective by mentioning the problem of the classical simulability of quantum states. For a meaningful quantum utility or advantage, the quantum state defined by the PQC must *not* be efficiently simulable classically. This remains an open research challenge, as it is unclear what ultimately makes a quantum state hard to simulate, either in the "weak" sense (efficient sampling of its squared amplitudes) or the "hard" sense (efficient computation of outcome probabilities). Note that, to simulate a VQA classically, it would be enough for the PQC to be simulable in the *weak* sense (i.e. sampled efficiently, to compute the cost function C and its gradients), a concept that has not been fully explored, yet [121]. Low-entangled states can often be simulated using tensor network techniques [122], while stabilizer circuits, relying only on Clifford gates, are efficiently simulable by the Gottesman-Knill theorem. Quantum entanglement and "quantum magic" are essential but insufficient on their own to ensure classical intractability. Neural network quantum states provide another powerful simulation method [72, 123].

In conclusion, we argue that integrating theoretical insights into application-focused studies is essential for advancing the pursuit of quantum utility. In particular, conducting large-scale simulations with shot noise while addressing BP mitigation represents a crucial step forward, alongside real hardware implementations with error mitigation techniques. Furthermore, the careful design of a PQC that is not classically simulable could pave the way for quantum speedups. At the same time, if a classically simulable PQC empirically outperforms conventional methods, it would lead to a powerful quantum-inspired classical algorithm. This may represent an equally valuable outcome in the broader landscape of quantum computing research.

3.3 Hybrid Quantum-Classical Neural Networks (HQC-NN) for RS and EO

In the previous section, we introduced the concept of a PQC trained in the VQA framework, mentioning that it is sometimes referred to as a Quantum Neural Network (QNN). Here, we further expand the terminology by introducing Hybrid Quantum-Classical Neural Networks (HQC-NNs). These models combine QNNs, made of parametrized quantum gates, with classical NNs, composed of standard artificial "neurons" arranged in different architectures, in particular variants of Convolutional Neural Networks (CNNs). Our interest stems from the observation that nearly all RS and EO applications discussed in the following sections utilize these hybrid architectures. In fact, limited size, noise, and error rates in current quantum hardware may severely limit the practical performance and reliability of fully-quantum NNs, whereas HQC-NNs represent an interesting alternative approach.

A summary of the concepts and tools of modern machine learning and Neural Networks is beyond the scope of our work and can be found e.g. in [124]. Nonetheless, App. B provides a concise overview of CNNs, which are central to the majority of reviewed papers, due to their ability to efficiently extract and learn hierarchical features directly from raw images. In essence, this is accomplished by applying interleaved trainable filters within *convolutional layers* and *pooling layers* for dimensionality reduction.

In Sect. 3.3.1 we briefly describe some common architectural choices for HQC-NNs, in particular quantum versions of convolutional and pooling layers that can be implemented *within* the quantum part of the network, namely in the PQC. We then review the published literature on RS and EO. We notice that HQC-NNs have been implemented in various ways, depending on the specific task assigned to the quantum subroutine.

The models discussed in Sect. 3.3.2 leverage the quantum layers of the network to perform the final feature transformation or classification task, while delegating data preprocessing and feature extraction to classical computational resources. Hence, the quantum part of the HQC-NN is placed at the end of the neural network and constitutes the final layers, often including the last one (output layer).

On the other hand, the works discussed in Sect. 3.3.3 are characterized by the use of quantum layers and quantum effects in the initial part of the network to perform feature extraction and dimensionality reduction from the input data. These quantum-enhanced features are then processed classically in subsequent layers with standard neural network architectures.

3.3.1 *HQC-NNs*: *theoretical overview*

The majority of proposed and tested architectures utilize CNNs and fully connected layers for the classical component of the network. Moreover, the concept of convolutional and pooling layers has been adapted into a quantum framework in various ways, resulting in multiple definitions of Quantum Convolutional Neural Networks (Q-CNNs) across the literature, in particular in Refs. [125–127]. The potential for confusion is further amplified by the inconsistent use of acronyms across the literature for Hybrid Quantum-Classical Neural Networks (HQC-NN), Quantum Convolutional Neural Networks (Q-CNNs), and Hybrid Quantum-Classical Convolutional Neural Networks (HQC-CNNs). We hope to clarify this issue by using a consistent notation, separating "quantumness" from the architectural details with clear hyphenation in acronyms.

The next paragraphs offer a concise summary of the main Q-CNN architectures, which have also been adapted to remote sensing applications. A clear understanding of these proposals could foster novel research in QML for RS and EO, e.g. by extending proof-of-principle tests to realistic benchmarks as discussed in Sect. 3.2.2.

Quantum convolutional neural networks by Cong et al. (2019) In the formulation of Ref. [125] from 2019, Q-CNNs are a particular version of a QNN with specific quantum convolutional and pooling layers inspired by their classical counterparts. This architecture is depicted in Fig. 6. In a Q-CNN, the input data is initially mapped onto a quantum state through quantum data encoding, or it is already quantum in nature. It is then transformed by a sequence of parametrized quantum gates. The i-th convolutional layer applies a *trainable* unitary operation U_i, which is quasilocal – i.e., it acts only on a few qubits – and is applied in a translational invariant way to the whole quantum register. On the other hand, the pooling operation in the j-th pooling layer is achieved through measurements performed on a fraction of qubits, whose results condition unitary rotations V_j on nearby qubits. Therefore, the pooling layers concurrently achieve nonlinearity through measurements and reduction in the number of qubits. It is also possible to replace intermediate

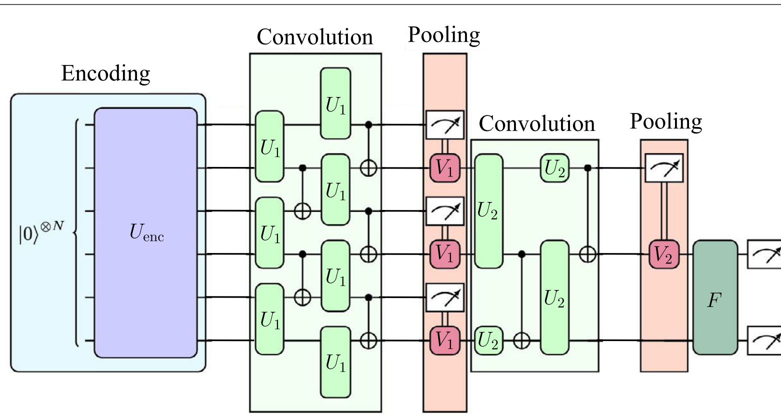

Figure 6 The main components of a Q-CNN applied to a learning task on classical data. This is an example of the general architecture explained in the main text. Figure adapted from Ref. [128]

measurements with unitary controlled rotations and perform all measurements at the end of the process [125]. The convolutional and pooling process is repeated a number of times until the system size is sufficiently reduced. A final fully connected layer transforms the remaining qubits by applying a trainable unitary operator F, and the resulting state is ultimately measured. Remarkably, only $\mathcal{O}(\log(N))$ variational parameters are used for input sizes of N qubits, simplifying the training process and reducing the computational costs for potential implementation on gate-based hardware. The training can be performed with a hybrid quantum-classical loop as described in Sect. 3.2.1. Typical loss functions involve the mean squared error (MSE) between the desired output and the Q-CNN output for given training quantum states:

$$\mathrm{MSE} = \frac{1}{2M} \sum_k \left(y_k - f_{\{U_i, V_j, F\}}(|\psi_k\rangle) \right)^2 \tag{22}$$

where y_k are the training labels, $|\psi_k\rangle$ are the training states, and $f_{\{U_i, V_j, F\}}$ represents the Q-CNN output function dependent on the unitary transformations. This formulation is a specific example of the general cost function in Eq. (15), where the variational parameters $\boldsymbol{\theta}$ are hidden inside the trainable unitaries. Other common choices include the cross-entropy loss function between the desired and the predicted labels. The original paper [125] utilized quantum input data and demonstrated impressive capabilities for quantum phase recognition and quantum error correction. Promising results for quantum phase detection have also been obtained in Ref. [128] by training only on marginal points of the phase diagram and proving the generalization capabilities of the Q-CNN to the whole phase diagram. However, these topics are beyond the scope of our discussion, which only focuses on architectural proposals and their application to QML for classical data.

Quanvolutional neural networks by Henderson et al. (2020) Ref. [126] from 2020 proposes a completely distinct scheme dubbed *quanvolutional neural network*. Differently from the fully-quantum model introduced in [125] and described above, this is an inherently hybrid quantum-classical architecture. The idea is to start from a classical CNN and replace one or more convolutional layers with new *quanvolutional* layers. Similarly to classical convolutional layers, a quanvolutional layer is composed of a variable number of filters, each responsible for extracting a different feature by acting locally on the previous layer. However, a given filter is realized by a *random* quantum circuit: first, the classical data is *encoded* in the quantum circuit, then it is processed by applying a random selection of one-qubit and two-qubit gates and, finally, the quantum state is measured to perform a *decoding* step that leads to the feature value, which is the input to subsequent classical layers. Since the random quantum state acts *locally* on the input data from the previous layer (it has a receptive field, see App. B), the number of qubits is small and tunable. It is crucial to notice that, here, the quantum circuit is *not* trained. Hence, there is no need for a quantum-classical optimization loop. On the contrary, it acts as a feature extractor by using a fixed and random quantum circuit: this leads to intrinsic noise resilience to real hardware noise if the error models are unknown but constant in time. Indeed, the specific choice and fine-tuning of gates is not relevant, and gate errors do not alter the proposed scheme.

A simple proof-of-principle test in the original publication shows interesting results on an MNIST benchmark, with a performance improvement over the standard CNN. However, the same results are obtained upon substituting the random quantum circuit with fully classical random noise models. Hence, no quantum advantage is present in this simple test. In follow-up works, several architectural hyperparameters could be tested. Starting from any given classical CNN, one should set the number and position of quanvolutional layers replacing convolutional layers, the number of quantum filters within each quanvolutional layer, the structure or randomness of the quantum circuits, and the size of each input qubit register. Additionally, one should define how to encode classical data into the quantum circuit and decode or extract the information from the quantum circuit, which is propagated to subsequent layers. A paper [129] that refines the original proposal and addresses EO datasets is discussed below in Sect. 3.3.3.

Quantum convolutional neural network for classical data classification by Hur et al. (2022) The authors of Ref. [127] test a Q-CNN scheme tailored for classification problems on classical data. The architecture is equivalent to the fully-quantum PQC introduced by Ref. [125], with an additional data encoding layer to embed classical data into the quantum circuit.

Despite the conceptual similarity to the work in [125], the focus of Ref. [127] lies in systematically investigating different Q-CNN implementations for classical datasets. In particular, the MNIST and Fashion-MNIST are selected as a benchmark. Several variants of Q-CNN models are tested and compared. Firstly, a number of different quantum data encoding techniques are compared, and secondly, different realizations of the trainable quantum convolutional circuit are tested. The authors also implement distinct classical preprocessing strategies, cost functions $C(\boldsymbol{\theta})$ (MSE and cross-entropy), and optimization algorithms. Remarkably, results on these standard testbed datasets numerically show better performance of the Q-CNNs compared to fully classical CNNs.

Although we do not cover the details of this comparison, we refer the interested reader to the original paper. In fact, this publication is particularly relevant for practitioners in QML for space science and technology, who are typically dealing with *classical* datasets, and it may provide several hints for designing quantum convolutional layers or fully-quantum CNNs. Additionally, a clear description of the training procedure in a VQA framework is provided, equivalent to the one in Sect. 3.2.1. Moreover, Appendix A of that paper draws a comparison among different versions of quantum and hybrid CNNs, including some proposals that seem incompatible with NISQ technology and are not described in our review.

As the authors highlight, key directions for future research lie in real-hardware implementations and their application to more realistic datasets. This presents promising opportunities also for advancing QML in the domains of EO and RS. A first step in this direction has been taken by Ref. [130], as discussed below in Sect. 3.3.2.

We conclude this section with a conceptual remark on Q-CNN for classical datasets. The Q-CNN architectures used in Refs. [125, 127, 130] closely resemble a multiscale entanglement renormalization ansatz (MERA) [131], with an additional fully connected unitary operator F acting on a final small subset of qubits. Due to their structural similarity to MERA, QCNNs can often be simulated classically using tensor network techniques, in particular when acting on a trivial initial state such as a blank register (or any state with low entanglement). Hence, the hope for the quantum utility of this proposal depends on the embedding procedure (or the input quantum state), which must be classically hard. Realistic simulations with shot noise and quantum hardware implementations, particularly for non-trivial classical datasets with classically intractable embedding schemes, will be essential to assessing the practical effectiveness of QCNNs for QML on classical data.

3.3.2 Applications of HQC-NNs to RS and EO: quantum layers for classification

The five works discussed below in Refs. [130, 132–135] share a common focus: they all apply QML in the form of hybrid quantum-classical schemes and leverage quantum layers to perform the final stages of feature transformation and/or classification on classical data for Earth Observation. A pictorial description of this framework is represented in Fig. 7, taken from Ref. [135]. In each case, the dataset consists of multispectral satellite images from the Sentinel-2 mission by the European Space Agency (ESA). Further technical details on the mission are summarized in an informative paragraph in Ref. [134].

Binary classification with HQC-NNs The works in Refs. [132, 133] tackle a binary classification task on the EuroSAT dataset [136]. This dataset contains Sentinel-2 data featuring 13 spectral bands and includes a total of 27,000 labeled, georeferenced images, categorized into 10 classes. To simplify the task, only three bands, typically RGB, have been selected, and the classification is performed with respect to two classes only.

In the first work, Zaidenberg et al. [132] compare the performance of an HQC-NN with that of a classical CNN. The classical architecture is composed of initial convolutional layers followed by fully connected (FC) layers and softmax classification. In detail, three convolutional 2D layers are followed by a ReLU layer, a max pooling 2D layer, and finally, an FC layer. This first part of the network architecture is also identically incorporated in the HQC-NN. After these steps, the classical version simply requires additional FC layers. In the hybrid quantum-classical case, however, the weight parameters of the first FC layer are used as rotational parameters of single-qubit Pauli Y rotations applied to an initial quantum state. This state is then measured, and the expectation values of computational basis measurements are interpreted as the weights connecting to the second and last FC hidden layer. This architecture is a variant of the standard LeNet-5 or AlexNet. Then,

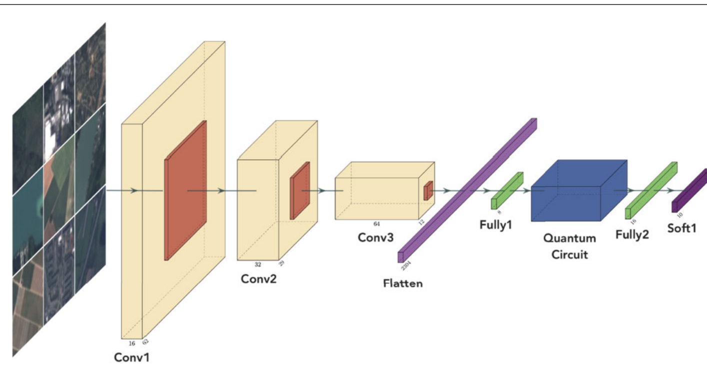

Figure 7 The input image is first heavily preprocessed with classical resources, often CNNs. The quantum layers of the HQC-NN perform the final stages of feature transformation and/or classification. In this specific example the quantum layers are inserted in between two fully connected classical layers. Figure from Ref. [135]

the final fully-connected layer of this network is arbitrarily weighted and mapped to the output layer.

The training of the model is done through backpropagation, where the parametrized angles of rotation in the quantum layer and the classical weights are optimized. The networks are used to perform several binary classification tasks of the RGB images of the EuroSAT dataset. The results show that the quantum neural network performs comparably or even better than the classical CNN, slightly improving its accuracy by 5 or 10% for the most challenging classification pairs, at the expense of an increased computational cost due to the additional quantum and classical FC layers in the HQC-NN. On top of the standard limitations due to decoherence and current shortcomings of quantum hardware, here, the actual utilization of quantum resources is very limited since single-qubit rotations can be simulated classically and do not create additional entanglement on the initial state. Nonetheless, this preliminary investigation paved the way for successive developments discussed below [135].

In the second work [133], Otgonbaatar and Datcu propose an HQC-NN where the classical part is made by two components. The first is a VGG16, which is a deep convolutional neural network, extracting meaningful features from the input data. It is followed by an autoencoder, specifically an energy-based model defined by a restricted Boltzmann machine (RBM), which performs dimensionality reduction of the input image. The quantum part of the network, consisting of a PQC, is applied subsequently. It has been designed and implemented on a classic simulator using the Tensorflow Quantum Python kit and is devoted to performing the binary classification task.

The EuroSAT dataset is used for training, while a real-world image of Berlin (Germany) from Sentinel-2 is used for testing. Each image in the dataset has dimension $64 \times 64 \times 3$, and needs to be reduced, through the autoencoder, to size $4 \times 4 \times 1$ in order to provide this classical input to the PQC, which has only 17 qubits. This limitation nearly matched the hardware state-of-the-art of that time and allowed for simple classical simulation with limited computational resources. One of the qubits is utilized as a readout qubit to encode the prescribed label value 0 or 1. Interestingly, the PQC is defined as an energy-based two-local model with trainable XX and ZZ couplings between the readout qubit and all the other "data" qubits, which encode the classically preprocessed features. The results show that a PQC classifies the EuroSAT images with the same high accuracy as a classic deep learning method, which is obtained by replacing the PQC with an FC classical layer. The HQC-NN happens to perform better for the complex remotely sensed image used as test data. In this work, the low number of qubits in the PQC is identified as the main limiting factor, as high-resolution images or high-dimensional data from more challenging datasets may not be reduced meaningfully to only 16 features.

Binary classification with Q-CNNs The binary classification task on EuroSAT, along with the SAT4 dataset, is also addressed in the work by Chang et al. in Ref. [130] by using a fully quantum convolutional neural network (Q-CNN) instead of hybrid quantum-classical architectures. Nonetheless, preliminary dimensionality reduction and feature extraction are performed with a classical convolutional autoencoder on the input data, and the classically preprocessed features are embedded into the QCNN. This makes it, in essence, an alternative implementation of the HQC-NN scheme using the quantum layers for classification. This work builds on the previous results by Hur. et al [127], and the authors compare different binary quantum classifiers. The architecture of the quantum network, inspired by Ref. [127], is composed of convolutional layers and pooling layers. The convolutional layers are variational quantum circuits on two qubits parametrized by a set of rotation angles for single- and two-qubit operations. The pooling layers reduce two-qubit states into a single-qubit state. Finally, a reduced single qubit in the output layer is measured to obtain the probability of class 0 or 1 membership. The parameters of the Q-CNN are updated to minimize the binary cross-entropy loss. The implementation in [130] differs from the original one introduced in [127] for a technical aspect. In fact, the latter uses the same variational parameters for all filters to ensure translational invariance; the former instead uses more parameters, different for all filters, breaking translational invariance. The preliminary results of this work show that this new model may perform better than the original one. This is somehow expected, as the number of variational parameters is significantly larger, also leading to higher computational costs. Two different quantum embedding techniques are compared. One uses Dense Qubit Encoding (DQE), and the other uses Hybrid Angle Encoding (HAE). We refer to the original paper for details. The preliminary results of this work show that the DQE embedding technique outperforms the HAE for most of the cases considered.

Multi-class classification with HQC-NNs Binary classification is a simpler version of the more general multi-class classification. The other two works cited above [134, 135] tackle this problem. The first seminal paper [134] is from 2020 and utilizes a PQC to perform Land Cover classification of a dataset created in the framework of the ESA "Sentinel-2 Global Land Cover" (S2GLC) project. Initially, the data undergo dimensionality reduction through principal component analysis (PCA), where only the most significant features are retained so that input data is projected onto a subspace equal in dimension to the number of qubits used in the quantum classifier. These features are then normalized and encoded into a quantum state as single-qubit Pauli X rotational angles. This is the first part of the so-called variational quantum classifier [109], which takes care of encoding the classically preprocessed features into the quantum circuit. The second part of the circuit, instead, consists of entangling variational layers that can be trained in the framework of a VQA with a hybrid quantum-classical loop. The variational circuit includes a sequence of entangling CNOT gates and trainable parametrized single-qubit rotations, a version of the so-called hardware-efficient ansatz [137].

The variational quantum classifier performs *binary* classification between any two specific classes by minimizing the mean square error between true labels $y = \pm 1$ and predicted labels. The latter are obtained by computing the expectation value of the total magnetization $\sum_j \hat{\sigma}_j^z$ on the full PQC, summing to it an additional trainable bias term, and then applying the sign function. The PQC variational parameters are iteratively adjusted to minimize the mean square error. Multi-class classification is done classically, by using the *OneVSOneClassifier* built on top of the binary classifiers obtained by applying the variational quantum classifier to each pair of classes. This approach achieves an overall accuracy of approximately 66% in multi-class land cover classification, below the accuracy levels achieved by state-of-the-art classical systems, which can exceed 86% accuracy. This suggests that further advancements in the PQC choice and computational scheme are necessary to match the best classical performance. For instance, using a PQC that is physically motivated or encodes some symmetries of the input data may prove beneficial [138].

Different kinds of data-encoding schemes, e.g. based on data re-uploading, could also be tested [139].

The work in [135] also addresses multi-class classification and is a development stemming from Ref. [132]. The goal here is to perform classification of the EuroSAT images for all the 10 classes available in the dataset. An HQC-NN was implemented as a modified version of LeNet-5 with two FC layers stacked before and after a quantum layer. The scheme is the same as described above for Ref. [132], where the features from the previous FC layer are used to determine the rotational angles of the quantum layer's PQC, whereas the computational basis (Pauli Z) expectation values on the quantum state determine the weights of the successive FC layer. However, here the quantum layer has been created and evaluated in three different ways by implementing three versions of PQC with four qubits. The first quantum circuit has no entanglement and applies first a Hadamard gate and then a parametrized Pauli Y rotation gate to all qubits. The second and third PQCs are defined as the *Bellman circuit* and *Real Amplitudes Circuit* respectively, and create entanglement between the qubits by applying Hadamard gates, parametrized Pauli Y rotations, and CNOT gates in different combinations. Contrarily to Ref. [134], this HQC-NN approach performs *direct* multi-class classification. Remarkably, the entangled models perform better both with respect to the other quantum model and to their classical counterparts. Moreover, the *Real Amplitudes Circuit* reaches the best accuracy of 92%, with a 10% gain over the *Bellman circuit*. This small-size experiment may suggest that entanglement can represent a valuable resource leading to increased computational capabilities. However, four qubits can easily be classically simulated, and larger-scale simulations on classical devices and experiments on real quantum hardware are mandatory for any solid claim of empirical quantum utility. This would entail addressing significant challenges such as barren plateaus and local minima, discussed in Sect. 3.2.

Recent advances towards practical utility The early works in QML for EO, including Refs. [130, 132–135] summarized above, have recently been extended in several ways. The main practical objective is to move from proof-of-principle results to quantum utility [28], namely the adoption of hybrid quantum-classical schemes as a valid alternative to classical methods.

A common challenge highlighted in the previous articles is *dimensionality reduction*, namely the necessity of downsizing satellite dataset images to adapt them to the typical input domain of PQCs utilized for classification. This is a key point for quantum computing in the NISQ era, since the number of qubits available on physical hardware is limited, and satellite imagery has too large spatial resolution and too many bands to be directly classified by quantum circuits.

A recent paper by Zollner et al. [140] systematically compares the performance of various dimensionality reduction techniques as part of hybrid quantum-classical systems. These dimensionality reduction techniques are used to represent satellite images from two benchmark datasets, EuroSAT and NWPURESISC45, performing a classical preprocessing of the input data, downsized from up to $256 \times 256 \times 3$ values to 16 values. The dimensionality reduction techniques are compared based on their impact on the performance of different versions of the variational quantum classifier [109]. The quantum classifier implementations test two different data-encoding schemes, namely basis and angle encoding, and four distinct trainable PQCs selected from standard literature: the

Farhi quantum classifier [141], the circuit-centric (CC) quantum classifier [109], the tree tensor network (TTN) quantum classifier [142], and the multiscale entanglement renormalization ansatz (MERA) quantum classifier [131, 143]. The theoretical motivation and physical interpretation of these circuits can be found in the original references. For practical purposes, we mention that the Farhi architecture has parametric Ising gates with trainable XX and ZZ rotations, whereas the other architectures include standard parametric single-qubit rotations and CNOT gates. The hybrid systems compared are, in total, 72, combining nine dimensionality reduction techniques for the classical pre-processing, two data encodings, and four PQCs for the quantum classifiers. They are compared by evaluating their performance on two binary classification tasks for each of the two datasets used. The nine dimensionality reduction techniques are discussed in detail in the original Ref. [140]. For instance, PCA and two distinct versions of an autoencoder (one of which is based on an RBM) are tested, with and without a prior feature extraction carried out by VGG16, the same deep convolutional architecture also used in Ref. [133]. The results show that autoencoders outperform other commonly used dimensionality reduction techniques, such as PCA, and may enable even small quantum circuits to have potential use for real-world applications. Concerning the comparison among the PQC architectures, the model proposed by Farhi et al. [141] shows the best performance among the ones considered in this work. This PQC is the same as the one used in Ref. [133].

As noted in [144], the use of a pre-trained deep convolutional NN for extracting features to encode into a quantum circuit is an implementation of quantum *transfer learning*. The preliminary results discussed so far appear promising, despite the intrinsic challenges of remotely detected datasets and the criticalities that can be encountered when applying a transfer learning scheme, as briefly mentioned in Sect. 3.1. The authors of Ref. [144] explicitly recognize this approach as a transfer learning scheme and adopt the same VGG16 architecture pre-trained on the ImageNet dataset. The PQC performing feature transformation and classification is either a real-amplitude or a strongly entangling circuit, with and without data re-uploading. During the PQC training, the classical network weights are kept frozen. The main incremental advance of this article lies in the dataset, which is characterized by a *small* number of images that are *high-dimensional*, i.e. not effectively represented with a small number of features. In contrast, the EuroSAT dataset is *big* and *low-dimensional*. The quantum transfer learning scheme is successfully applied to a three-class dataset, where the classes are *dense residential*, *medium residential*, and *sparse residential* areas taken from the UC Merced Land Use dataset. Only three qubits are used, and a connection between performance and a theoretical measure known as *local effective dimension* is tested empirically. Additional numerical work on hard data sets is mandatory to empirically test HQC-NNs for remote sensing. However, a larger number of qubits should be used under more realistic experimental conditions (including shot noise or real hardware implementations).

We finally point out the *recent technical paper* in Ref. [145], which could be useful for practitioners interested in this field of application. Here, the authors do not propose new architectural advancements but rather adopt a side perspective and focus on other topics. They perform a comparison of different quantum computing libraries routinely used to implement VQAs, out of which PennyLane may be the most versatile due to its integration with PyTorch GPU libraries. In addition, they assess the relevance of the initialization values chosen for the variational parameters of a PQC by evaluating their impact on the final performance. Finally, they propose integrating quantum layers into Vision Transformer (ViT) classical architectures. Once more, the dataset is composed of RGB bands of multispectral images from the EuroSAT dataset. The authors also provide an interesting historical excursus of a selected subset of papers on quantum and hybrid approaches for EO and RS.

3.3.3 *Applications of HQC-NNs to RS and EO: quantum feature extraction*

The works discussed in this section approach hybrid quantum-classical schemes from a different perspective than those in Sect. 3.3.2. Here, the quantum components of the network are utilized at the initial stages of the workflow, as schematically depicted in Fig. 8. Specifically, quantum layers are positioned early in the network to handle feature extraction and dimensionality reduction. The benchmark problem is still the supervised learning and classification of multispectral EO datasets.

Quantum layers for image representation and feature extraction In the previous section, quantum computing has primarily been employed to transform classically extracted features or to directly perform the classification step. However, a conceptual limitation of this approach lies in its computational costs: indeed, feature extraction and pre-processing are typically more resource-demanding than the final classification process. In practice, the classical subroutines utilized in the HQC-NN described above often use computationally intensive deep convolutional neural networks for image comprehension and feature extraction. In contrast, only a simple PQC is typically used for the final steps, often with a very limited number of qubits.

The authors of Ref. [147] propose to invert the two steps and use a complex PQC to extract important features from input images and a simple classical FC layer to perform the final classification. We remark that real-world data, particularly EO images, are large and multidimensional, and the encoding into quantum states is challenging given the small number of qubits available in current NISQ devices or efficiently simulable in classical emulators. Ref. [147] uses a technique named FRQI (Flexible Representation of Quantum Images) [148] encoding an image of size $2^N \times 2^N$ by using $2N + 1$ qubits. The position of pixel j (e.g. in raw-major order) is represented by the computational basis state j of $2N$ qubits, whereas the additional qubit is devoted to encoding the pixel intensity (e.g., grayscale value) as a rotation angle. This quantum state, in general, requires an exponential number of quantum gates (controlled rotations) to be prepared. This limits the practical applicability of FRQI only to moderate image sizes. Compression schemes have been

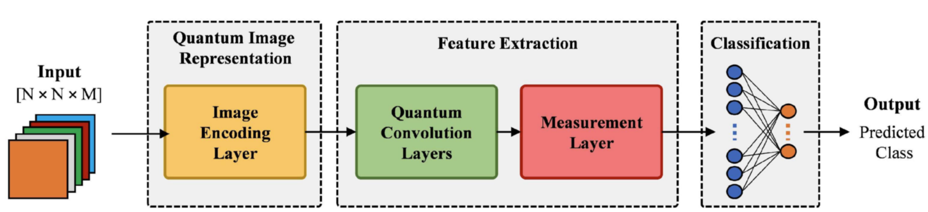

Figure 8 In this algorithmic pipeline, the quantum layers of the HQC-NN perform feature extraction, after the preliminary encoding of an $N \times N$ image with M spectral bands into the quantum circuit. The final feature processing and classification steps are handled by classical layers. Figure adapted from Ref. [146]

devised [149], and are utilized in Ref. [147]. The second part of the PQC implements quantum convolutional layers, which extract high-level features in a way conceptually similar to convolutional layers in classical CNNs, as outlined in Sect. B. Implementation details of the convolutional architecture are described in the original reference, but the main components are multi-controlled single qubit rotations, where the controls are the position and value of a given pixel. The third part of the PQC is made of the measurement layer, which translates the quantum features extracted by the quantum convolutional layers to classical feature values. The quantum state is measured in the X-basis, Y-basis, and Z-basis, to obtain the corresponding expectation values. The latter are concatenated and serve as input to the fourth and final part, which is a *classical* dense layer to perform the final classification. This preliminary work tackles multi-class classification on a simplified version of the Overhead MNIST dataset, composed of gray-scale images (a single color band). The HQC-NN model achieves performance comparable to that of classical CNN for this highly simplified problem.

This approach has been extended and improved in the recent publication [146] by the same authors. In this work, they propose two different models of HQC-NN to perform land cover multi-class classification of Sentinel-2 multispectral images. Both hybrid quantum-classical methods exploit quantum computation to perform encoding of the images into the quantum register and feature extraction, whereas classical computation is used for the final classification task. These models leverage quantum convolutional layers, and could thereby be named Hybrid Quantum Classical Convolutional Neural Networks (HQC-CNN). The first model uses MCQI (Multi-Channel Quantum Images) to perform image encoding into the quantum register and can be dubbed (in our notation) M-HQC-CNN. The second model is a direct generalization of the one implemented in the previous work [147], using FRQI, and dubbed F-HQC-CNN. Both models perform the multi-class classification of multispectral images through three steps: quantum image representation, feature extraction, and classification.

In the M-HQC-CNN model, the first step exploits the MCQI method, which is suited for encoding images with multiple bands on a single quantum circuit, by using a well-defined number of qubits to specify the pixel location, the band index, and the pixel value. On the other end, the F-HQC-CNN model exploits the FRQI method, which was originally proposed to encode classical gray-scale images and uses multiple quantum circuits, one for each spectral band of the image. Multiple-controlled qubit rotations are the essential building blocks of both methods, and a detailed explanation is given in the original reference. For the second step of feature extraction, the quantum circuit ansatz is rooted in the same concepts described in the previous publication [147]. Essentially, quantum convolutional layers extract high-level features, and a measurement layer converts quantum features to classical values via the computation of specific expectation values obtained from Pauli measurements. The quantum ansatz is slightly different in the two quantum-classical models: in the F-HQC-CNN model, multiple quantum convolutional circuits are needed for feature extraction from each band separately, and therefore, a distinct measurement layer is used for each band, instead of a single one as in the M-HQC-CNN model. As noted by the authors, the M-HQC-CNN has the potential to select more refined features by combining information from different bands. However, this comes at the expense of using more qubits in the quantum circuit to encode the band index. For the final classification task, both models use a *classical* dense layer to classify the classical feature vector output from the measurement layer. Standard training of the variational parameters in the quantum convolutional layers is performed using the cross-entropy loss function.

Ref. [147] addressed the LCZ42 dataset, a benchmark dataset for global Local Climate Zones Classification, with some technical simplifications. The image size was reduced from 32×32 to 8×8 pixels, only four bands out of ten were selected, and only the images labeled with three German cities were used, instead of all 42 cities.

To assess the performance of the two models proposed in this work, these were compared with a classical CNN with two convolution layers, each comprising 2 and 24 kernels, respectively, and with the HQC-NN discussed above for multi-class classification, introduced in Ref. [135], where a classical CNN was operated as a feature extractor. A fair comparison is performed by choosing a similar number of trainable parameters for the competitor models.

The two models using quantum features, i.e., F-HQC-CNN and M-HQC-CNN, lead to better average test accuracy and superior transferability and robustness. The benchmark has been extended to other models for binary classification using only the RGB bands. We refer the interested reader to the original paper for a detailed statistical analysis and insightful discussion. Overall, quantum layers for feature extraction show promising results and may represent a valid paradigm for future research. Crucial steps forward include simulating a larger number of qubits with at least shot noise instead of exact state-vector simulations, and, ultimately, tests on real quantum hardware. Additionally, tackling hyperspectral datasets with many bands would be an interesting way to assess the potential of quantum features extracted by mixing different band information, e.g. by adapting the M-HQC-CNN framework.

Application of the quanvolutional architecture In Ref. [129], the original proposal of quanvolutional NNs [126] has been extended and tested for the classification of low-resolution satellite imagery, specifically a subset of the SAT-4 dataset. These are $28 \times 28 \times 4$ images (namely, with four bands) with 4 distinct labels: barren, trees, grassland, or other.

In this work, the authors use a single quanvolutional layer in the first part of the architecture; hence, they test quantum layers for *feature extraction*. Nonetheless, we remark that the quanvolutional NN framework is flexible and the quanvolutional layers may be placed early or at the final stages of the network, or even in more places at the same time. This version of quanvolutional NN has been tested on real hardware with a Rigetti computer.

The original architecture introduced in [126] has been refined significantly. Firstly, the quantum circuit is not random but rather structured; in particular, the authors choose a QAOA variational circuit (with depth $P = 1$, which is the main QAOA hyperparameter [34]). Furthermore, the encoding and decoding schemes are improved and made realistic, relying on measurements rather than a full state-vector representation of the quantum state in a quanvolutional filter.

Referring to the *quanvolutional* architecture described in Sect. 3.3.1, the authors in [129] use a single quanvolutional layer of 5-by-5 quanvolutional filters, requiring 25 qubits. The quanvolutional layer is applied to $5 \times 5 \times 4$ patches of the original images, corresponding to a total of 100 feature components to be encoded. The encoding scheme is an angle encoding into the QAOA variational angles that parameterize X rotations. Since these are 25 per QAOA layer, in principle, one should choose depth $P = 4$. To reduce the circuit depth, the authors rather consider $P = 1$ and encode in each angle an average over 2×2 sub-patches of contiguous pixels.

Interestingly, a software and hardware co-design approach has been implemented by selecting a QAOA circuit requiring two-qubit gates that are natively implemented in the hardware topology of the Rigetti quantum device. The paper does not provide sufficient details regarding the algorithmic implementation, particularly concerning the interpretation of the ZZ rotation angles and the decoding scheme. Based on the available information, we hypothesize that these angles are initialized randomly for each quantum filter and kept fixed, coherently with the original proposal of Ref. [126]. Assuming this is correct, QAOA is *not trained* here but rather adopted as an *encoding circuit*, by generalizing the angle encoding scheme in Eq. (18).

The comparison with a classical CNN does not show any clear advantage, yet this formulation could be improved and expanded in many ways, along the same lines discussed above for the original proposal [126].

3.3.4 Other space science and technology applications of HQC-NNs

Quantum-classical computational schemes, in particular implementing variants of hybrid neural networks, have been investigated for a wide range of application areas, including space science and technology beyond RS and EO. A comprehensive review of the whole literature would be impractical. Nonetheless, we briefly mention a few relevant papers that apply these techniques to interesting space applications. Our scope here is dual. First, we point the interested reader, with a technical background in space-related sectors, to potentially interesting literature. Secondly, we highlight the conceptual similarity of different hybrid quantum-classical proposals, remarking that the computational scheme and the main conceptual and technical challenges are often unrelated to the application domain.

A first interesting example in Ref. [150] concerns *space weather detection*. Here, an HQC-NN is proposed to enhance the detection and early warning systems for space weather events, such as solar flares and geomagnetic storms, which significantly impact space technologies and infrastructures. The authors devise an HQC-NN that achieves promising performance. The quantum layer with six qubits is positioned between two classical dense layers and uses angle embedding encoding, entanglement, a quantum activation function, and measurement operations. Remarkably, the experiments were also tested with real hardware on IBM quantum services.

In the study conducted in Ref. [151], the application domain is *forecasting solar irradiance*, a task traditionally addressed by deep learning techniques. This research investigates the integration of quantum layers into conventional feedforward neural networks and the development of a fully quantum PQC for the learning task, without any quantum convolutional layer. To benchmark these advancements, a bidirectional long short-term memory (BiLSTM) model is employed. The quantum-enhanced feedforward neural networks display robust and competitive performance, despite using significantly less variational parameters. As expected, the fully quantum PQC model exhibited sub-optimal performance in comparison, probably due to its too simple architecture.

Ref. [152] introduces an HQC-NN architecture tailored for near real-time *space-to-ground communication* of International Space Station (ISS) data, detected with the Lightning Imaging Sensor (LIS). In the long term, such hybrid architectures might be promising for processing and transmitting vast amounts of sensor data, which classical systems struggle to handle efficiently. Initially, a classical network handles feature extraction from the ISS LIS lightning dataset. These features are then encoded into quantum states using multiple quantum encoding techniques. Performance comparisons between classical models and the HQC-NN seem promising in terms of training and validation accuracy.

An interesting approach that falls into the category of RS is explored in Ref. [153]. Here, QML is applied to *PolSAR (Polarimetric Synthetic Aperture Radar)*, an advanced form of SAR (briefly introduced in Sect. 3.1) that exploits the polarization of electromagnetic waves to provide additional information about the Earth's surface, e.g. for land use and land cover applications. Here, the authors map the polarization state of the reflected beam (reflected wave's Jones parameters) into single-qubit states, a procedure dubbed *natural* embedding, which exploits the natural equivalence between the Poincaré sphere for light polarization and the Bloch sphere of a qubit. They also test an HQC-NN architecture to classify this data, with the quantum layers applied to extract quantum features, before the classical ones perform the classification task. This follows a similar workflow as the one described in Sect. 3.3.3.

Finally, we mention Ref. [154], which tackles a complex binary classification task for an EO dataset, namely the classification of satellite images to determine whether they contain solar panels or not. This paper uses classical pre-processing and develops two distinct quantum models based on *neural quantum kernels* (NQKs). This technical tool is discussed in detail in the original publication and shows promising results.

3.4 Quantum Support Vector Machines (QSVM) for RS and EO

Support Vector Machines (SVMs) are non-parametric supervised machine learning algorithms that can be employed for classification and regression problems. Compared to deep neural networks or other machine learning models, SVMs are particularly suitable when the training set is small and the acquisition of additional training data is expensive or impossible. Another key feature of SVMs is their stability, namely, small alterations of the training data do not generally produce significant differences in the trained model.

SVMs are particularly appealing for RS applications, where typically only limited training samples are available, and further acquisitions would require impractical economic and time overheads. An extensive amount of research has been devoted to SVMs in the RS field, from urban growth monitoring to forest detection, analyzing different types of images with a different number of bands, and with resolutions ranging from submeter to many kilometers for each pixel [155].

Recently, a number of proposals to devise a Quantum Support Vector Machine (QSVM) have been put forward. After a concise overview of classical SVMs in Sect. 3.4.1, we review the proposal of a QA-based SVM and its applications to RS, in Sects. 3.4.2 and 3.4.3, respectively. The same is done for gate-based SVMs in Sects. 3.4.4 and 3.4.5. Our discussion focuses on the main conceptual novelties and common challenges that are inherently present in these hybrid quantum-classical ML algorithms.

We remark that the QA-based SVM does not use gate-based quantum computation and is inherently an application of Quantum Annealing, as illustrated in Fig. 1. We have chosen to present it here to simplify the exposition. Conversely, gate-based SVMs utilize quantum kernels, which involve embedding classical data into a quantum circuit, as outlined in Sect. 3.2.1. These kernels may also incorporate variational parameters that are optimized within a hybrid quantum-classical loop.

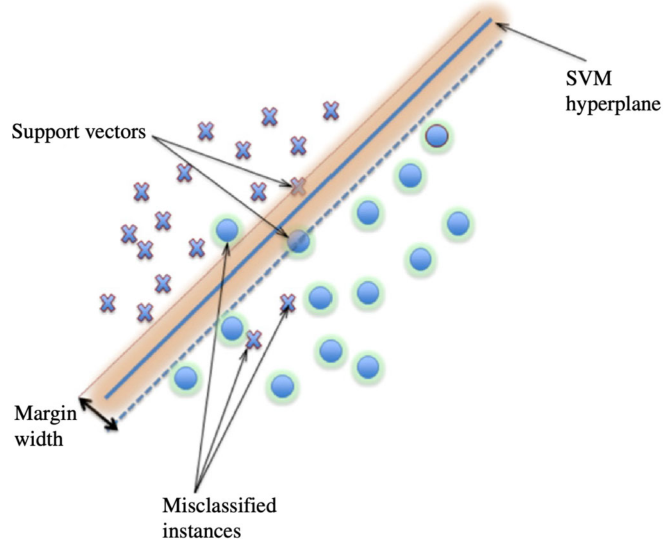

Figure 9 The geometric interpretation of a linear SVM, which is the simple formulation of Eq. (1) without soft margins or the kernel trick. Figure from Ref. [155]

3.4.1 A short overview of the theory of classical SVMs

An SVM operates on a labeled training set

$$\mathcal{D} = \{(\mathbf{x}_n, y_n) : n = 1, \dots, M\}, \tag{23}$$

where $\mathbf{x}_n \in \mathcal{X} \subset \mathbb{R}^d$ is a feature vector and $y_n = \pm 1$ its prescribed label. A test set is defined as $\mathcal{T}$, whose samples are also drawn from $\mathcal{X}$, with unknown labels that are statistically correlated with those of the training set.

Intuitively, an SVM separates the training samples of two distinct classes in the feature space by tracing a maximum margin hyperplane between them, as shown in Fig. 9. The hyperplane is defined by its normal vector $\mathbf{w}$ and offset b, which should be determined such that for any training point $\mathbf{x}_n$, we have

$$y_n(\mathbf{w} \cdot \mathbf{x}_n - b) \geq 1. \tag{1}$$

In order to minimize the risk of *test* data being misclassified, one also aims to maximize the margin width $\frac{2}{\|\mathbf{w}\|^2}$, geometrically representing the distance between the closest vectors belonging to different classes. The latter are called *support vectors* and have a crucial role in classifying test data after the training phase, as explained below.

This basic formulation of a linear SVM is generalized by including soft margins that allow for a few misclassified data points and the *kernel trick*, enabling the definition of non-linear decision boundaries separating the two classes [156]. Training an SVM is a convex quadratic problem subject to convex constraints; thus, it always converges to the global minimum. This implies that there is no need for repeating training using different random initializations or stochasticity. The mathematical theory behind SVMs is refined and can be found, e.g., in Ref. [156]. This theory leads to an equivalent dual formulation for SVM training, which is immediately relevant for the definition of QA-based QSVMs [157].

By considering the general case that includes kernels, training an SVM is equivalent to a quadratic programming problem [157]:

$$\text{minimize} \quad E = \frac{1}{2}\sum_{n,m=1}^{M} \alpha_n \alpha_m y_n y_m k(\mathbf{x}_n, \mathbf{x}_m) - \sum_{n=1}^{M} \alpha_n, \tag{24}$$

$$\text{subject to} \quad 0 \le \alpha_n \le C \quad \forall n = 1, \dots, M, \tag{25}$$

$$\text{and} \quad \sum_{n=1}^{M} \alpha_n y_n = 0, \tag{26}$$

where the M optimization variables are $\alpha_n \in \mathbb{R}$, while C is a regularization parameter and $k(\cdot,\cdot)$ is the kernel function.

The kernel trick relies on mapping the original d-dimensional vectors to a larger (and possibly infinite-dimensional) Hilbert space $\mathcal{H}$, whose elements are now called *feature* vectors. Although this is formally realized with a map $\Phi : \mathbb{R}^d \mapsto \mathcal{H}$, only scalar products of feature vectors are needed for training an SVM (see Eq. (24)), which define the kernel function:

$$K(\mathbf{x}_i, \mathbf{x}_j) = \Phi(\mathbf{x}_i)^\dagger \cdot \Phi(\mathbf{x}_j)\,. \tag{27}$$

The above equation defines the Gram matrix for the training set $\mathcal{D}$, which is Hermitian and positive semi-definite. The majority of classical kernels are defined in terms of real feature maps. A notable example is the radial basis function (RBF) Gaussian kernel, often just called the RBF kernel:

$$K(\mathbf{x}_i, \mathbf{x}_j) = e^{-\gamma \|\mathbf{x}_i - \mathbf{x}_j\|^2}\,, \tag{28}$$

where γ is a tunable parameter, and the norm is the conventional Euclidean 2-norm. Here, $\mathcal{H}$ is infinite-dimensional, and the explicit map Φ would be computationally intractable. The kernel trick avoids this problem by only computing kernel functions, which can be done easily, without requiring an explicit expression for the feature map.

A typical solution often contains many $\alpha_n = 0$. It can be proven that the only non-zero variables $\alpha_n \neq 0$ correspond to those feature vectors that lie on the decision boundary and are named the *support vectors* of the SVM. A prediction for an arbitrary test vector $\mathbf{x} \in \mathcal{T}$ can be made by computing the sign of the decision function:

$$f(\mathbf{x}) = \sum_{n=1}^{M_\mathrm{S}} \alpha_n y_n\, k(\mathbf{x}_n, \mathbf{x}) + b, \tag{29}$$

where the sum has been restricted to the M_S support vectors. An explicit expression for the bias b can be provided, or it can be set to zero by definition [157].

As summarized in the conclusions of Ref. [155], SVMs emerge as a competitive candidate for performing ML tasks for RS images due to their effectiveness for training sets with a limited size or low-quality training data. Intuitively, this is due to the definition of an optimal boundary separating the two classes depending only on the support vectors. In addition, there is no requirement for a prior assumption for the data distribution, which is unknown for most RS real-world use cases.

However, some limitations include the arbitrary choice of kernels, especially if the original data in $\mathbb{R}^d$ are already high-dimensional, and the selection of the regularization parameter C. Finally, while the SVM framework could, in principle, be generalized from binary to multiclass classification in different ways, several training issues can emerge.

3.4.2 QA-based QSVM: theory

As anticipated in Sect. 2.7, D-Wave quantum annealers have been utilized for training Support Vector Machines in various applications for the space sector. In fact, the training process of an SVM has been equivalently reformulated as a QUBO task in Ref. [157]. The main idea behind QA-based SVMs is to discretize the quadratic programming problem in Eq. (24), by replacing the float precision values of the variables $\{\alpha_n\}$ with a discrete set of values represented in some encoding base. Each variable is represented using K bits, ensuring that the constraints in Eq. (25) are automatically satisfied. The optimization is then recast to a QUBO problem with KM binary variables (or qubits), where a Lagrange multiplier is introduced to account for the constraint in Eq. (26). For a kernel as in Eq. (28), whose values decay rapidly as a function of the distance between two feature vectors, the couplings between several qubits of the QUBO graph are small and can be neglected, simplifying the embedding process.

As noted by the authors in [157], two remarkable conceptual advantages may emerge in this approximate formulation of the training of a classical SVM with a quantum annealer. First, determining the convergence to zero of a floating point number α_n can be subtle, whereas, with a discrete mesh of values including zero, it is straightforward. This is a crucial step in identifying the support vectors for the trained SVM. Secondly, the global minimum found by direct solution of the SVM quadratic programming problem is optimal only for a certain *training set*. On the other hand, quantum annealing samples may provide several close-to-optimal solutions corresponding to different decision boundaries in the feature space. A combination of these has the potential to outperform the classical global minimum when evaluated on the *test set*, which is the ultimate goal in supervised classification.

Due to the current size limitations of QPUs, the authors propose to train a set of QA-based SVMs only on a subset of the original training set, finally combining them to obtain an aggregated classifier. This approach is found to be competitive or even outperforming standard SVM training on a synthetic and a realistic datasets.

3.4.3 QA-based QSVM: applications to RS and EO

In this section, we summarize the key contributions to the application of QA-based SVMs in the field of remote sensing. The summary is presented in chronological order, emphasizing the incremental progress made in the field.

Binary classification of multispectral images The methods outlined in the previous section have found numerous applications to several datasets of practical interest in remote sensing and Earth Observation. In Ref [158], a set of QA-based SVMs are trained on disjoint subsets of RS multispectral images for binary classification. These are aggregated together to perform inference on the test data, showing encouraging preliminary results.

This work has been extended in Ref. [159], utilizing the newer Advantage D-Wave Quantum Processor. The authors compare QA-based SVMs to gate-based SVMs, which we shall introduce in the following section. The former approach, relying on the QUBO formulation described above, delivers better results. The authors consider a dataset dubbed SemCity Toulouse, which is a multispectral image dataset focused on the city of Toulouse. For our purposes, it offers a useful, practical example: here, an urban area is described as a multispectral image of $500 \times 500 = 250{,}000$ pixels with 8 bands. Only 50 pixels are selected for training, each consisting of a feature vector $\mathbf{x}_n \in \mathbb{R}^8$. The two labels $y_n = \pm 1$ correspond to the absence or presence of a building in that pixel.

Core of a dataset and weighted SVM An alternative strategy to apply the QA-based SVM to challenging EO datasets has been presented in Ref. [160]. The idea is to replace the original dataset with a "core of a dataset" (or coreset), which is a small representative weighted subset of the original training set. This coreset technique has been applied to train a weighted version of the standard SVM with D-Wave hardware. A preliminary tuning of D-Wave annealing parameters is performed on two simple coresets of a synthetic dataset and the Iris dataset. These parameters include the annealing time τ, the chain strength J_F, and the number of samples to be drawn (see Sect. 2.2). These values of the annealing parameters are thereby used to train the QA-based weighted SVM for binary classification on the coreset of two real EO datasets. Results on classification accuracy are competitive with a state-of-the-art classical weighted SVM.

Extensions to regression and multiclass classification The QA-based SVM framework has been extended to regression analysis for estimating biophysical variables in remote sensing [161]. Furthermore, a generalization to multiclass SVMs was proposed in Ref. [162]. The authors extend the original proposal of Ref. [157] to perform direct multiclass classification with an SVM trained with a quantum annealer. The QUBO mapping now leads to a matrix with a linear dimension equal to KMC, where K and M are defined above, and C is the number of classes of the training set. Two different multispectral RS datasets were considered. The classification accuracy is nearly comparable to the classical implementation of the same multiclass SVM. Interestingly, this quantum approach leads to a favorable scaling of the execution time as a function of the training set size M: in particular, the inference time is constant in M. Although the embedding process of the QUBO graph becomes increasingly challenging with M, this trade-off between the execution time and the size of the QUBO graph may be favorable in some applications.

3.4.4 Gate-based QSVM and quantum kernels: theory

A first proposal for a Quantum Support Vector Machine was introduced in the seminal paper [163], based on the inversion of a kernel matrix (introduced below) through an adaptation of the celebrated HHL algorithm [164]. However, this version of QSVM would require fault-tolerant quantum computation and is incompatible with current NISQ devices. Moreover, its theoretical exponential quantum speedup over classical SVM could only be obtained if classical data were provided in a quantum superposition, which is a highly idealized scenario; hence, we will not explore its formulation. We instead focus on a hybrid quantum-classical approach based on quantum feature maps that encode standard classical data into a quantum Hilbert space, which plays the role of a high-dimensional feature space [165, 166]. A gate-based quantum computer is used to compute quantum kernels, which can be utilized to train a *classical* SVM as described in the previous Sect. 3.4.1.

The underlying idea is to embed a classical data point $\mathbf{x} \in \mathbb{R}^d$ by defining a *quantum feature map* $\mathcal{U}_{\Phi(\mathbf{x})}$. This is implemented with a parametrized quantum circuit that acts on an initially blank register of N qubits in the state $|0\rangle^{\otimes N}$, preparing the state

$$|\Phi(\mathbf{x})\rangle = \mathcal{U}_{\Phi(\mathbf{x})} \, |0\rangle^{\otimes N} \, . \tag{30}$$

The quantum feature map is often defined in terms of a set of fixed entangling gates and a set of parametrized Pauli rotations, whose rotational angles are determined by the components of the input vector $\mathbf{x}$. This may be regarded as quantum data encoding (or embedding) as discussed in Sect. 3.2.1, often a generalization of angle encoding in Eq. (18). The selection of an appropriate PQC to embed classical data into a quantum Hilbert space is crucial. A valid approach relies on the introduction of additional variational parameters $\boldsymbol{\theta}$ in the quantum feature map, which can be tuned to optimally separate the training set in the feature space using, for example, a method known as kernel-target alignment [167]. The training of parametrized kernels is performed within a VQA scheme using a hybrid quantum-classical optimization loop.

To correctly formalize the mathematical description of quantum kernels within the framework of classical kernel theory, the correct formulation would be to represent quantum states as *density matrix operator*, which for a pure state reads $\rho(\mathbf{x}) = |\phi(\mathbf{x})\rangle\langle\phi(\mathbf{x})|$ and has the additional advantage of eliminating the unphysical global phase of ket states [168].

Looking back at Eqs. (24) and (29), a quantum computer would be needed to compute the kernel matrix between pairs of training samples (during training) or between a test sample and support vectors (during inference), i.e. to estimate

$$K(\mathbf{x}, \mathbf{x}') = \mathrm{Tr}\left[\rho(\mathbf{x}')\rho(\mathbf{x})\right] = |\langle\phi(\mathbf{x}')|\phi(\mathbf{x})\rangle|^2 = |\langle 0|^{\otimes N} \, \mathcal{U}^{\dagger}_{\Phi(\mathbf{x}')} \mathcal{U}_{\Phi(\mathbf{x})} \, |0\rangle^{\otimes N}|^2 \, . \tag{31}$$

Notice that, formally, a linear Hilbert-Schmidt scalar product is indeed defined in the density matrix representation, where we recall that any density matrix operator satisfies $\rho = \rho^{\dagger}$ (it is also positive semidefinite with unit trace). This formulation is the so-called fidelity quantum kernel, since its evaluation requires us to estimate the *global* quantity in Eq. (31), namely the fidelity between two pure quantum states obtained by quantum-feature encoding of two distinct classical data points.

Estimating fidelity on a quantum computer can be done via standard methods such as the SWAP test or the overlap test (sometimes called the Loschmidt Echo test). The latter consists of a direct evaluation of the probability of measuring all qubits in state $|0\rangle$ after preparing the ket state in the last member of Eq. (31), obtained by the subsequent action of a feature map circuit and an adjoint feature map. See e.g. Ref. [169] for additional details.

However, it is important to note that the fidelity between two distinct quantum states is expected to vanish exponentially fast with an increasing number of qubits, due to non-equal quantum states becoming orthogonal in the exponentially large Hilbert space. Designing a quantum feature map involves a delicate trade-off. On one hand, it must be inherently difficult to simulate classically to enable the possibility of quantum speed-up. On the other hand, recent studies on quantum kernel trainability have highlighted that exponential concentration (a phenomenon equivalent to barren plateaus [117]) may often occur. In summary, the kernel values $K(\mathbf{x},\mathbf{x}')$ may all concentrate exponentially fast in the number of qubits N around a mean value, independent of the input data $\mathbf{x}$ and $\mathbf{x}'$. For example, when evaluated using the overlap test, the quantum kernel Gram matrix would converge to the identity matrix. This results in a trivial model that becomes completely ineffective when a polynomial number of measurements is used to estimate its values. Exponential concentration may arise due to highly expressive or highly entangling quantum feature maps, but also noise and, notably, global observables such as the fidelity between two quantum states [169].

Alternative approaches, such as the projected quantum kernel [170], relax this strong requirement of estimating a global quantity. However, also this method is susceptible to trainability issues, especially related to entanglement in the quantum feature map, or hardware noise. We refer to [169] and references therein for recent developments on trainability and best practices in quantum kernels, such as exploiting symmetries or devising problem-inspired quantum feature maps. Finally, it is worth noting that quantum kernels can serve as an alternative to traditional kernels, such as the Gaussian kernel in Eq. (28), for a wide range of applications extending beyond the SVM framework.

3.4.5 Gate-based QSVM: applications to RS and EO

In this section, we outline a selection of the main applications of quantum kernels to remote sensing. The first paragraph describes an insightful analysis on cloud detection using the fidelity quantum kernel, which is simulated on a classical computer. We then move to land-use and land-cover classification with classically simulated projected quantum kernels. Finally, we focus on a seminal paper reporting a real implementation of a gate-based SVM on a noisy quantum processor, where a dataset for the binary classification of supernovae into types II and Ia was selected as a challenging use case.

We also cite Ref. [171], where the authors examine the issue of concentration in quantum kernels. They practically evaluate the trade-off between the exponential cost of classical simulations and the exponential number of measurements required – due to concentration – for reliable quantum kernel estimation on a quantum device. This reference also provides an introductory and high-level summary of the main QML challenges along with a partial list of applications to remote sensing.

Cloud detection in multispectral satellite images A thorough application of hybrid SVMs enhanced with quantum kernels is proposed in Ref. [172], where the authors approach the task of cloud detection in satellite imagery data, specifically analyzing a large-scale dataset of practical relevance composed of multispectral images from the Landsat-8 satellite.

The binary classification of image pixels into clear and cloudy areas is a crucial step in the field of EO for two main reasons: first, cloudy regions are often removed from the dataset because the target ground areas are likely to be obscured. Secondly, efficient cloud detection is instrumental for meteorological and climate research. Furthermore, the removal of cloudy areas represents a key step for on-board image compression and dataset reduction.

The 38-Cloud dataset contains a set of diverse landscape images, where each pixel is characterized by four spectral bands and a ground truth label indicating the presence or absence of clouds at that specific point of the image. A classical dimensionality reduction step is performed through a superpixel segmentation technique, followed by a classical feature extraction with principal component analysis (PCA).

The quantum feature map is implemented using the PennyLane Python package and consists of multiple layers: a data encoding layer S that depends on the input data point $\mathbf{x}$, a variational layer W, and an entangling variational layer E. Both S and W involve only single-qubit rotations and do not introduce entanglement. Notice that W and E include trainable variational parameters $\boldsymbol{\theta}$, which are trained to optimally embed the training set in the feature Hilbert space by utilizing kernel-target alignment [167]. This method is clearly and thoroughly summarized in [172], and we recommend referring to it for further technical details. The investigated architectures are implemented by sequentially stacking the previously defined blocks, specifically adopting the S, WS, ES, and $WSWS$ configurations.

For each architecture, quantum kernels are classically simulated and used to enhance the training of an SVM. Its performance is compared to a standard radial basis function (RBF) kernel SVM, which utilizes the kernel in Eq. (28) with an optimized parameter γ. In summary, the gate-based SVM offers a classification accuracy comparable to the classic radial basis function (RBF) kernel and superior to a trivial linear kernel, where the feature map Φ is the identity and therefore $K(\mathbf{x}_i, \mathbf{x}_j) = \mathbf{x}_i^\dagger \cdot x_j$.

Interestingly, entanglement is not necessary for this application, as the introduction of E layers does not improve the performance of the model. A quantum feature map without entanglement can be efficiently simulated classically with polynomial resources in the number of qubits N, excluding any potential quantum advantage. However, these results hold for a specific binary classification task performed on a heavily classically preprocessed dataset and should not be extrapolated to a general context. In fact, only 4 features were used, and in the embedding scheme implemented by the S layer, this is equivalent to setting the number of qubits to $N = 4$.

This work has the potential to be generalized to other real-world datasets in RS and EO, perhaps using more features and simulating more qubits, or by implementing this scheme on real quantum hardware.

Land-use and land-cover classification with projected quantum kernels Another crucial task in EO is to classify land use and land cover in image datasets collected from satellite missions. In Ref. [173], the authors perform supervised classification on multispectral images of the Earth surface obtained from the Copernicus Sentinel-2 mission. These images are characterized by 12 spectral bands and 16 land use labels. As discussed above, both classical and quantum-enhanced SVMs are more frequently employed for binary classification tasks. Similarly, this work also focuses on the simplified task of separately distinguishing between two pairs of land use labels.

The QML tool used in this paper is Projected Quantum Kernels (PQK), introduced and motivated as part of a long theoretical work published in [170]. The complex theory presented in this article has been applied to a specific use case in [173], deferring the description and motivation of many technical details to the original reference. Therefore, we also refer to [170] for a detailed exposition on PQKs. Concisely, the idea is to redefine the quantum kernel formalism by avoiding comparing two feature vectors in an exponentially large Hilbert space, as done for fidelity quantum kernels in Eq. (31). One way to do this is to define the PQK in terms of the one-particle reduced density matrix for each single qubit. Starting from the density-matrix operator representation of the quantum feature state for a data point $\mathbf{x_i}$, namely $\rho(\mathbf{x_i}) = |\phi(\mathbf{x_i})\rangle\langle\phi(\mathbf{x_i})|$, one obtains the one-particle reduced density matrix for qubit k by tracing out all other qubits indexed by $j \neq k$, leading to

$$\rho_k(\mathbf{x_i}) = \mathrm{Tr}_{j \neq k}\left[\rho(\mathbf{x_i})\right] . \tag{32}$$

The PQK is finally defined as

$$k^{PQ}(\mathbf{x_i}, \mathbf{x_j}) = \exp\left(-\gamma^{PQ} \sum_{k=1}^{N} \|\rho_k(\mathbf{x_i}) - \rho_k(\mathbf{x_j})\|_F^2\right), \tag{33}$$

where γ^{PQ} is a real positive hyperparameter and $\|\cdot\|_F$ is the simple Schatten 2-norm.

Once again, classical pre-processing is applied to the original dataset, including PCA leading to the selection of a subset of relevant features. The authors compare the performance of two identical classical feedforward neural networks with three layers, only one of which has access to features extracted from the input image using a PQK approach. They performed binary classification on a modified version of the dataset, where they re-assigned the labels to maximize the geometric distance (defined in [170]) between the PQK in Eq. (33) and an RBF classical kernel as in Eq. (28) with optimized hyperparameter γ. This empirical study demonstrates an improved classification accuracy on the test set for the model utilizing PQK features. Although still in its early stages, this work stands out from the reviewed literature by applying a sophisticated mathematical framework to a practical use case in EO. It could pave the way for further research in quantum kernels for remote sensing applications beyond the basic formulation of fidelity kernels.

Implementation on quantum hardware for supernovae classification A significant distinction between the reviewed literature on QA-based SVM in Sect. 3.4.2 and the study of gate-based SVM using quantum kernels lies in the implementation on real quantum hardware. In fact, all papers using QA have implemented their proposals on D-Wave hardware, whereas the quantum kernels are frequently simulated on classical devices rather than executed on actual quantum hardware. This represents a major shortcoming: not only is hardware coherent and decoherent noise fully ignored, but also the statistical "shot noise" due to the inherent measurement process in quantum mechanics is neglected by running full state-vector simulations. We remark that shot noise is ultimately related to trainability issues such as exponential cost concentration [169] discussed in Sect. 3.4.4, which would be easily avoided if one had access to a full-precision state vector for the quantum state.

A remarkable exception is offered by Ref. [174], presenting a full-stack 17-qubit implementation of quantum kernels on the Google Sycamore processor based on superconducting qubits. Interestingly, their methods were tested on an RS dataset, which makes the discussion in this paper particularly well suited to the scope of our work. The dataset is composed of astronomical time series with flux measurements across six wavelength bands. The authors focus on binary classification between Type II and Type Ia supernovae. Only minor preprocessing has been applied to produce 67 features, which are embedded in a *real* quantum circuit composed of 17 active qubits without further dimensionality reduction.

Contrarily to the previously cited work on cloud detection [172], in Ref. [174] the authors choose a non-parametric quantum kernel method. The quantum feature map is selected to preserve, in practice, kernel magnitudes of non-diagonal Gram matrix elements, which would otherwise vanish due to exponential concentration. Tailored error mitigation techniques have been devised to improve the performance of noisy simulations on the Sycamore quantum processor.

For a given data point, each feature is embedded using angle encoding, represented as the rotation angle of a Pauli gate applied to a specific qubit. It is important to note that the number of qubits (17) does not necessarily need to match the number of features (67) in this encoding scheme.

Rigorous statistical analysis is performed, leading to the remarkable conclusion that noise-mitigated real-hardware simulations provide similar classification accuracy compared to ideal noiseless simulations performed on a classical computer. However, hardware simulations are characterized by significantly larger training errors. As the authors suggest, hardware noise seems to prevent the model from overfitting the training data, a phenomenon that needs to be further investigated, since it could be beneficial for inference on test data (namely, for classification accuracy).

4 Integrating quantum computing, communication, and sensing in satellite missions: a perspective

In this section, we outline recent advancements and prospects of integrating quantum computing with other quantum technologies for space science and technology. As anticipated in Sect. 1, we do not intend to fully cover this extensive interdisciplinary field, which has already been thoroughly reviewed in prior studies [7–9]. Instead, we focus on satellite missions, specifically on the perspective of combining quantum computing with quantum communication and quantum sensing. A high-level representation is shown in Fig. 10, which illustrates a schematic configuration of a modular quantum satellite capable of acquiring, processing, and transmitting quantum information, and connected to a network of equivalent devices. This pictorial representation reflects the content of the next sections.

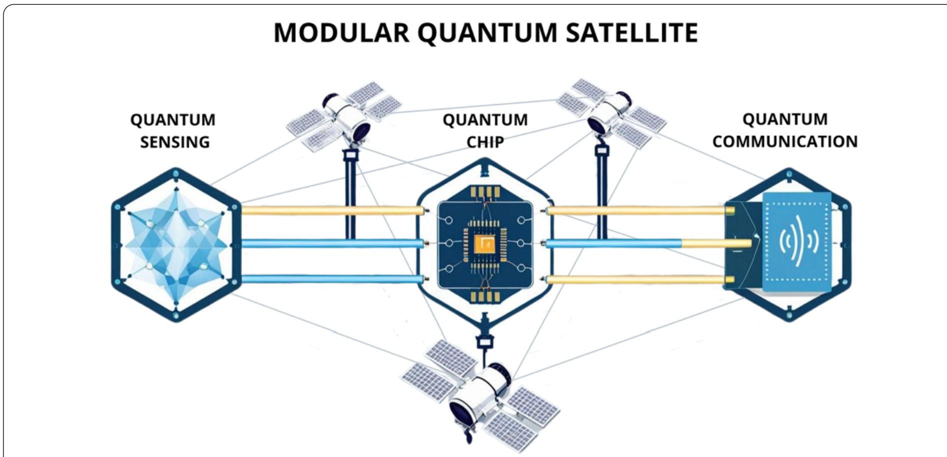

Figure 10 Graphical representation of a quantum satellite, consisting of three interconnected modules. Each module performs a distinct function and supports applications in a specific area of quantum technology, integrating quantum sensing, quantum computing, and quantum communication. The satellite is also connected to other similar devices, forming a constellation. The long-term goal is to devise such a network of quantum satellites capable of exchanging flying qubits and processing them on solid-state platforms. In particular, we consider nitrogen-vacancy (NV) centers in diamond, as illustrated by the leftmost module (see Sect. 4.2)

We begin in Sect. 4.1 by exploring the long-term objective of establishing a distributed quantum computing platform in space, leveraging a satellite constellation to integrate various devices implementing different quantum technologies. The first experimental implementations have demonstrated quantum communication networks, featuring both ground-to-space communication and the exchange of quantum information between satellites [175]. Based on recent advancements, we review the actual components needed for a fully operational quantum network.

Section 4.2 examines the promising application of quantum sensing deployed on satellite constellations, potentially leveraging quantum processors for computational tasks. We first focus on spaceborne data acquisition and Nitrogen-Vacancy (NV) centers in diamonds, and then discuss the potential of quantum sensing for multi-messenger astronomy and fundamental tests in physics.

4.1 Quantum networking and its space deployment

The first implementations of quantum networks have been conceived and tested to execute communication protocols. We refer to [176] for a discussion of the underlying physics and a thorough analysis of the challenges posed by the realization of a quantum communication network that integrates quantum memories. We also point out the supporting roadmap discussed in Ref. [177], where the authors have critically examined the potential benefits and performance of a satellite-based quantum information network across various sectors of activity. The authors in [178] offer a compelling survey on space quantum communication, addressing physical constraints such as atmospheric attenuation and their impact on achieving significant milestones over the past decades. Additionally, they provide valuable insights into feasibility studies exploring innovative research directions.

As mentioned in Sect. 1.1 and described in App. A, quantum computing and quantum technologies are now at a crucial stage, with various experimental approaches competing to become the standard for qubit technology. Early advances were achieved with superconducting qubits and trapped ions, while neutral atoms have recently gained ground, achieving the first implementations of logical error-corrected qubits. This accomplishment may represent a shift from the NISQ era towards error-corrected logical qubits [179]. Besides these three technologies, researchers also investigate spin qubits in silicon and photonic qubits. Photonic qubits, in particular, present a unique hybrid approach, combining elements of both solid-state and optical technologies, which is expected to play a crucial role in the development of a quantum internet.

Indeed, one of the primary challenges lies in reliably interconnecting quantum devices and generating entanglement across long distances. Terrestrial optical fiber networks, despite their potential for quantum communication, face prohibitive exponential losses, particularly over intercontinental distances, whereas deploying numerous repeaters introduces impractical costs and complexity [180]. It has been proposed that a quantum satellite internet constellation could serve as a promising alternative by mitigating signal attenuation and enabling global-scale quantum communications [181–183]. A pictorial representation is given in Fig. 11.

In Sect. 4.1.1, we broaden our discussion to explore the prospects of networks for distributed quantum computing, whereas Sect. 4.1.2 examines the conceptual building blocks needed to realize a quantum network, with a focus on foreseeable satellite implementations. Finally, Sect. 4.1.3 reports some key experiments of quantum communication that have been proposed for or realized in space.

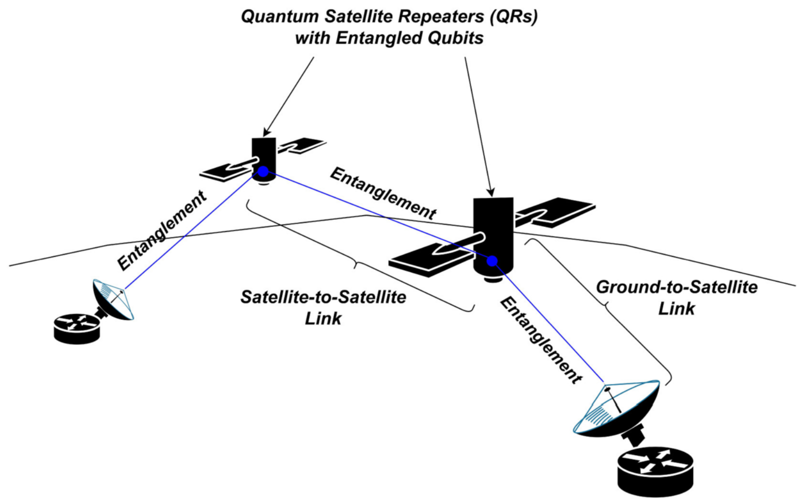

Figure 11 Pictorial description of the main constituents of non-terrestrial quantum networking from Ref. [184], to which we refer the reader for a detailed explanation

4.1.1 Prospects of a quantum network for distributed quantum computing

Moving beyond quantum communication networks, the ultimate goal would be to establish a quantum network for quantum communication and computation, building a global quantum internet that could enable distributed quantum computing and quantum computing in cloud [185]. The framework required for a quantum satellite internet constellation would share a similar conceptual architecture as a quantum network of computing units implementing distributed quantum computing on Earth. The primary conceptual challenges lie in storing and retrieving quantum information [176]. An effective solution of the latter would enable a distributed network of entangled QPUs to be realized on the ground via optical-fibers links. The conceptualization of such distributed architecture would not be altered by considering ground-to-satellite and satellite-to-satellite links. The extension to satellite nodes could be rather swift, thanks to the ongoing efforts in the quantum communication sector [184]. Nevertheless, as we write this review paper, the experimental implementations of distributed quantum computing schemes are still in their infancy [186].

An architecture implementing distributed quantum computing tasks over a network of computing nodes would allow the end-user to virtually access a number of qubits that scales linearly with the number of interconnected processors, diminished by a small fraction of qubits devoted to the communication tasks (see [187] for a detailed analysis). In the near term, the objective is to implement the distributed quantum computing paradigm with the computational ability of NISQ and beyond-NISQ devices before fault-tolerant quantum computing. Recently, Ref. [188] has proposed an approximated distributed computing scheme to implement variational quantum algorithms and time evolution of quantum spin glasses model on already available quantum simulators.

Similarly to classical distributed computing, which was enabled by prior technological achievements such as the development of ARPANET (the first network of communicating computers), a crucial aspect relies on the interplay of computation and communication in a distributed quantum computing environment [189]. In the quantum setting, this balance is particularly critical when qubits need to interact across potentially vast distances, and it is expected that optimal performance may be achieved using solid-state technologies for computation [190] and photons for communication [191]. In fact, photonic qubits are the ideal candidates for *flying* qubits, namely those used for communication, as they can travel over long distances either on optical fiber or in free space, with the possibility of exploiting different wavelengths (typically in the visible domain or the telecom band) to reduce losses, according to the experimental conditions. A paradigmatic architecture can employ spins in solid-state [192] or NV centers [193, 194] in diamonds that emit or absorb photons in the telecom spectrum, facilitating inter-module communication without requiring the physical proximity of qubits for computation. The connection of qubits across different quantum processing unit nodes with photonic links could potentially be scaled, reaching a higher computing power than isolated processors (see e.g. Fig. 1 in Ref. [187]).

The ultimate goal would be to develop a grid of quantum computers, potentially even NISQ devices that are not error-corrected, by letting them exchange quantum information and data in such a way they could act as a single machine. This approach also highlights the essential role of entanglement distribution. In contrast to classical communication and computing networks, a quantum implementation relies on the certification of a certain amount of quantum resources to function. In particular, entanglement allows quantum information to be shared between the quantum units, and improvements in entanglement distribution are expected to enable more scalable and efficient distributed quantum computing [195], as it is done, for example, by achieving noise reduction in Ref. [196]. In this context, quantum logic gates operating on different nodes to create maximally entangled states of two qubits have been demonstrated [197]. Additionally, the implementation of teleported controlled-NOT (CNOT) gates has been reported [198–200], providing compelling evidence that qubits in spatially separated quantum processors can be entangled. Moreover, a proof-of-principle experimental demonstration of distributed quantum computing, in which the implementation of the Grover search algorithm on a an architecture involving two trapped-ion modules interface via a photonic link has been reported in Ref. [201].

4.1.2 Building blocks of a quantum network

Here, we outline the building blocks of a grid of quantum units deployable in space that could serve as the backbone infrastructure of an extended quantum network. Such architecture could support several tasks ranging from quantum communication to quantum sensing for fundamental studies and practical applications. We now sketch the key concepts and challenges for the design and implementation of the long-term goal of a quantum internet, as explained in greater detail in Ref. [202]. The objective is to build a network of quantum devices that either directly interact or exchange quantum states. Linked multiple devices would enable the secure transmission of quantum information across distances and could find early applications in quantum cryptography, such as quantum key distribution and blind quantum computing [203].

From a structural perspective, a quantum network is largely similar to its classical counterpart but is characterized by its exploitation of quantum resources. The essential building blocks of a quantum network include quantum nodes, quantum channels, quantum repeaters, and protocols for quantum communication. Reliable quantum information manipulation and transfer critically depend on these constituents, which can be implemented across a variety of platforms. As a result, they represent a primary endeavor of current research in the quantum technology community.

Quantum networks are based on *quantum nodes*, which serve as processing or storage centers for quantum information. Such nodes often consist of three units accountable for acquiring the information, storing, and processing it, respectively. They can also feature a batch of quantum sensors, a slot for quantum memory, and, in the long term, they may include also a quantum computer. At the present stage, there are several proposals for the physical realization of quantum nodes. Common platforms include photonic chips, trapped ions or atoms, superconducting qubits, and color centers in diamond [185, 204], each of them with unique properties in terms of coherence time, operational temperature, and quantum error rates. For example, trapped-ion qubits have high coherence times but are challenging to scale, while superconducting qubits are more scalable but tend to have shorter coherence times [205, 206]. Moreover, the latter interact with cosmic rays [207], potentially leading to decoherence and dissipation, which are responsible for the loss of the quantum resources that are necessary for any quantum protocol. The physical implementation of a quantum node directly influences the performance and robustness of the network.

Quantum channels provide the link that allows transferring quantum information encoded into quantum states among different quantum nodes. The feasibility of exchanging photons across large space distances for quantum technology purposes was demonstrated soon after the first tabletop experiments in quantum information science [208, 209], with the Italian National Space Agency playing a pivotal role. The aim of such experiments is the characterization of the channels that allow the transmission of flying qubits. This can be achieved either through a physical medium or in a vacuum, and according to the physical nature of the carriers, several types of losses, dissipation, and decoherence effects can come into play. For optical communication, these channels can be fiber optic cables or free-space links. When deploying a quantum network in space, photon-based qubits are currently the most suitable information carriers, primarily due to the resistance of photons to decoherence in a vacuum and their ability to traverse free-space paths with minimal loss [210]. Photonic qubits have proven valuable in enabling ground-space quantum communication, particularly at wavelengths around 1550–1560 nm. These wavelengths offer high air transparency for free-space transmission and are well-suited for long-distance optical fibers [211] once on the ground.

The most detrimental effect for optical fiber channels on the ground is photon losses. Recent studies have shown robust techniques to lower the losses over long metropolitan distances. A notable example is offered by Ref. [212], using high-dimensional photonic entanglement [213] over 10.2 km with high noise due to sunlight exposition. The authors show that encoding the information on high-dimensional degrees of freedom can effectively preserve quantum correlations despite photon loss, enabling stable quantum key distribution protocols under daylight conditions. In Ref. [214], the authors tackle the challenge of overcoming the rate-loss limit in long-distance point-to-point quantum key distribution (QKD) experiments. They demonstrate secure key rates exceeding the standard over 413 km and 508 km of optical fibers. A key innovation in their approach is the elimination of the need for a common global phase reference, namely a shared phase alignment between communicating parties typically required for many QKD protocols. This eliminates a significant experimental constraint, making the system more practical and robust for real-world implementation. In a broader perspective, quantum channels underpin the crucial task of entanglement distribution across the network, and more advanced techniques, such as wavelength division multiplexing, are being investigated to find the optimal values for efficiency and capacity of quantum channels [215].

Quantum repeaters are a critical element in overcoming the distance limitations of quantum channels [216]. As already mentioned, due to photon losses and decoherence, quantum signals degrade over long distances, creating the need for a mechanism to transfer entanglement across wide networks. Quantum repeaters can be understood as intermediary quantum nodes with storage and entanglement-swapping abilities. Entanglement swapping is a process that allows two distant objects, initially in a separable state, to become entangled without direct interaction. This is achieved by first generating entanglement between intermediate, adjacent nodes and then performing a specific measurement operation on these intermediate nodes. This measurement transfers (or "swaps") the entanglement to the distant nodes, thereby establishing a long-distance entangled state. This approach enables the formation of an extended entangled link. Its discovery was of paramount importance to achieve quantum teleportation [217].

By breaking a long channel into shorter segments, quantum repeaters can generate entanglement locally and perform entanglement swapping to extend it further. This approach theoretically allows for an entangled state to be distributed over distances far exceeding what would be achievable in a direct transmission [218]. In practice, quantum repeaters require highly reliable quantum memory and controlled entanglement swapping processes, which are currently challenging to implement. Recent advancements in solid-state quantum memories, such as those based on rare-earth ions and atomic ensembles, offer promising avenues toward effective quantum repeaters [219]. We also highlight that the application of space-borne quantum memories is anticipated to have a disruptive potential across several research lines in space science and for fundamental tests from space [220]. The experimental scheme outlined above, where information is elaborated and stored in solid-state processing units and transmitted with photons, presents several technological challenges. A step forward is reported in Ref. [221], where the authors report a deterministic protocol to store and retrieve light from a quantum dot in an atomic ensemble quantum memory at telecom wavelengths.

Besides the hardware components, the development of a distributed quantum computing system deployed in space also requires significant advancements in *software stacks*. A fully functional constellation of satellites executing quantum operations would require two main software components. The first one, which may be either classical or quantum, is a software-defined networking approach designed for the effective management of available resources. This component should coordinate the entire constellation and may be operated from Earth or a station that actively manages the generation of entanglement among the nodes [184, 222]. Secondly, a distributed quantum computing paradigm necessitates an efficient compiler that matches all the physical constraints imposed by the different computing nodes. A crucial task is the optimization of resources needed for local operations and classical communications, and the more computationally expensive protocols for entanglement generation and distribution [187, 223].

4.1.3 Key experiments of quantum communication in space

At the present stage, the most successful protocols implemented on prototypical quantum networks are the cryptographic ones based either on the BB84 [224] or the E91 [225] protocols. Beyond the implementation of Quantum Key Distribution (QKD), which is the most mature application, other early proposals include teleportation-based and distributed quantum computing protocols that have been developed to facilitate information transfer and share computational tasks across quantum networks [226]. A fully operational quantum network for distributed quantum computing would ideally require the synchronization of quantum operations and error correction to manage decoherence and maintain fidelity across the network. All the hardware and software elements described in the previous section should be effectively integrated [184].

Despite the significant technological challenges, experimental progress in recent years has brought us closer to the realization of proof-of-concept quantum networks in space. Here, we highlight a few notable examples. One significant advancement is the development of a quantum communication station within the CubeSat standard, which consists of cube-shaped satellites with compact dimensions and lightweight design. These satellites typically have side lengths ranging from 0.1 to 1 meter and weight of the order of a few kilograms. They have been employed in the Q^3Sat experiment, in which a quantum-secure key was generated [227] using a minimal configuration. Additionally, their application has been proposed for missions deployed from the International Space Station [228], offering a flexible platform for quantum communication experiments. Such applications could pave the way for the launch of a dedicated Quantum Science Space Station, as discussed in [229].

Finally, we mention the experiments performed by the Chinese National Space Agency subsequent to the launch of a 600 kg satellite carrying out the Quantum Experiments at Space Scale (QUESS) project. Among its notable achievements, the transmission of entangled photons over record distances stands out, further demonstrating the viability of satellite-based quantum communication [175].

4.2 Quantum processing units for quantum sensing in space

We now shift our focus to the integration of quantum sensing with quantum processing units that may be achieved in the foreseeable future. Similarly to quantum computing processors, quantum sensors require a physical platform capable of preserving quantum properties for a prolonged time. However, in contrast to quantum computing, quantum sensors present fewer challenges, as they can be implemented even with a single qubit and do not necessarily require large-scale entangled systems [230]. The technological advancements and potential integration of quantum sensing and computing, combined with progress in controlling quantum systems both on the ground and in space, offer promising perspectives. The theoretical framework of VQAs [41] suitable to run on NISQ devices has recently been extended to implement experimentally programmable quantum sensors [231], or to use trainable quantum algorithms for estimation and metrology tasks [232–234].

Section 4.2.1 presents a synthetic description of applications of quantum sensing for data acquisition and remote sensing (RS) from space, along with a specific focus on Nitrogen-Vacancy (NV) centers in diamonds. In Sect. 4.2.2, we outline the potential of quantum sensing for multi-messenger astronomy and fundamental physics tests.

4.2.1 Quantum sensing for spaceborne RS and data acquisition

Quantum sensing has emerged as a promising approach for spaceborne Earth Observation (EO) [97], which is crucial for understanding Earth's physical and biological processes along with tracking changes in land use and land cover, as discussed in Sect. 3.1. Quantum sensing may lead to significant improvements over traditional EO instruments, potentially offering higher accuracy, sensitivity, and stability [235]. In view of the successful development and ground-based testing of various quantum sensors – some of which are commercially available – their application in space holds a broad potential [7] but poses significant challenges.

From a wider perspective, the deployment of quantum sensors in space is one of the most promising technologies for improving the quality of navigation systems [236] and for spaceborne data acquisition in general. Several platforms to implement quantum technologies in space have been proposed or even tested depending on the specific application. A complete summary goes beyond the scope of this section, but we mention a few relevant examples. Cold atoms are employed as atomic clocks allowing for precise geodesy and quantum gravimetry [237] measurements, whereas Bose-Einstein condensates are utilized for fundamental studies in a microgravity environment [238]. Other sensors are implemented with different technologies that could support deployment on an orbiting station. Among them, we cite quantum receivers based on the excitation of neutral atoms to high energy levels (also known as Rydberg atoms) [239, 240] and qubit sensors implemented in NV centers in diamonds [241]. Both technologies are particularly advantageous due to their high sensitivity to minute variations in external fields, such as electromagnetic fields, as well as their ultra-wide bandwidth and broad sensitivity range. Additionally, quantum LIDAR (Light Detection and Ranging) sensors can be used to detect nonclassical optical states of lights for extracting classical parameters (physical, measurable properties described by classical physics) of target objects [242].

A system should satisfy a few prerequisites to be considered suitable to operate as a quantum sensor. Crucially, it should possess a discrete series of energy levels, with the ability to isolate and define a pair of levels as a two-level system. The initialization of the system in its ground state should be easy and require minimal engineering effort, and the read-out state has to be achievable through coherent manipulation, e.g. via time-dependent electromagnetic fields [230]. One of the most promising implementations for emerging quantum sensor technologies relies on Nitrogen-Vacancy (NV) centers in diamonds [192, 243]. Indeed, NV centers naturally serve as qubit systems and can be envisioned as integrable components within a qubit-based quantum computer, forming part of a more complex device. This leads to the potential use of quantum sensors for acquiring data that can be subsequently directly transmitted to quantum hardware for processing. Among many advantages, this approach may reduce or even eliminate the need to store classical data and, consequently, the challenge of determining the optimal encoding of this data into a quantum algorithm. For most quantum sensors currently available, we do not yet see a clear path to developing such an operating device that relies solely on quantum components. On the other hand, it appears feasible to engineer a fully quantum data acquisition and processing system using NV centers by integrating the approaches for communication [244] and computation [21].

NV centers in diamonds represent a quantum system where a nitrogen atom substitutes a carbon atom in the diamond lattice, adjacent to a vacancy. The spin degree of freedom of this system acts as an extremely sensitive *compass needle*, capable of detecting minute variations in magnetic fields with high precision. The NV center operates by being excited with nonresonant green laser light and microwave signals. The green laser light excites the system, causing it to emit red fluorescent light. By applying microwaves across different frequencies and monitoring how this red fluorescence changes, the local magnetic field strength can be accurately measured. These systems can operate at room temperature and potentially withstand harsh conditions, making them ideal candidates for space applications. Although NV centers exhibit notable resilience and stability properties, a key challenge to address lies in optimizing the sensor properties to address the complexities of space missions. These include sensitivity to temperature fluctuations, noise levels, and interference from the hosting satellite.

The compactness of NV centers, along with their ability to work individually or collectively, offers the potential for high-resolution magnetic field mapping. The potential of NV center technology, especially its ability to make precise magnetic field measurements, is being investigated and tested in space [235]. Applications include mapping Earth's magnetic field and monitoring space weather using satellite constellations [243]. Moreover, magnetometry offers an invaluable tool to probe the interiors of planetary bodies, enabling remote investigation of a planet's internal composition, structure, and even historical evolution by analyzing its magnetic history preserved in crustal rocks. This technique can also reveal hidden subsurface oceans. NV centers are currently being employed in space mission research in the context of planetary-scale magnetic field measurements by leading space agencies.

Due to their sensitivity to small variations of external fields, the properties of NV center sensors are also suitable for magnetic navigation [245], in conditions where the global positioning systems are not available [246]. Their implementation as navigation systems deployed in satellite constellations with embedded magnetometers is currently under study [247, 248].

4.2.2 Networks of quantum sensors for multi-messenger astronomy

Multi-messenger astronomy explores astronomical phenomena and physical processes by integrating observations from various *messengers*, namely by different carriers of physical information originating from cosmic sources. Quantum sensors, such as atomic clocks, magnetometers, and gravitational wave detectors, hold great potential in this field, specifically for advancing fundamental physics. A promising perspective lies in the realization of quantum sensors that could operate in a network configuration and leverage quantum computing stacks to analyze the acquired data [249].

A key challenge lies in employing these sensors in the quest for exotic physics beyond the standard model. For instance, these networks can be repurposed to detect macroscopic dark matter objects or signals from exotic low-mass particle fields [250], particularly by capturing bursts generated by highly energetic astrophysical events. This can be enabled by the high level of accuracy achieved in shielding quantum processing units from conventional noise sources, such as electromagnetic radiation, cosmic rays, and other environmental disturbances. Another detection strategy uses the matched-filter technique to identify transient dark matter signals, such as objects moving at galactic velocities, and apply this to simulated data from GPS atomic clocks. The analysis in Ref. [251] demonstrates a sensitivity improvement of up to three orders of magnitude over previous GPS-based results, showcasing the broader applicability of this methodology.

Lastly, we highlight a promising application achievable with a network of computing units properly engineered to acquire and elaborate quantum data. Connecting multiple optical telescopes through a network of quantum computing nodes could dramatically improve the imaging of faint astronomical objects. This approach would involve encoding the quantum state of incoming photons along with their arrival times into qubit-based codes at each receiver. By using entanglement-assisted parity checks, the network can retrieve the overall quantum state nonlocally, by avoiding common transmission losses between nodes. This strategy preserves and extracts the crucial phase differences needed for interferometry, as outlined in Ref. [252]. Such methods could be integrated with quantum estimation schemes to go beyond the classical resolution limit, heuristically summarized by Rayleigh's criterion. According to this criterion, two incoherent point sources are only resolvable if they are separated by a distance proportional to the wavelength of the detected light multiplied by the ratio of the distance to the sources to the pupil size. Incoherent and faint emissions are recurrent in astronomical observations to discriminate the constituents of multiple systems, such as binary or trinary stars, or for the detection of an exoplanet orbiting around a much brighter star. An experimentally achievable scheme to achieve sub-Rayleigh precision in the detection of close incoherent sources has been theoretically proposed in Ref. [253]. The authors have derived the ultimate quantum limit in the form of the quantum Cramèr-Rao bound and proposed a realistic measurement scheme, known as spatial-mode demultiplexing. This method is based on the decomposition of the incoming light into Hermite-Gaussian spatial modes and avoids measuring the position of each incoming photon. Such a scheme only requires the implementation of linear photonics and has been experimentally demonstrated in [254]. Moreover, quantum metrology and state discrimination techniques, such as quantum hypothesis testing, have shown an advantage over classical methods, particularly in improving the scaling of the Fisher information and the Chernoff bound. A significant application arises in scenarios that can be formulated as binary decision problems, with exoplanet science being a promising area. In Ref. [255], the authors proposed a method for distinguishing a faint source emitting near a much brighter one, achieving a lower error probability compared to direct imaging techniques.

5 Outlook

In this review, we screened and analyzed published work on quantum computing for space science and technology. We adopted the perspective of quantum computing experts, organizing the presentation by sorting previous works by quantum algorithms and architectures utilized to solve a computational problem.

We aimed to offer an accessible and pedagogical summary of the key theoretical and algorithmic challenges in quantum computing, which are often independent of the specific application domain. In fact, a significant part of our effort is devoted to providing a timely, concise, and in-depth overview of quantum computing in the quest for quantum utility. We covered both Quantum Annealing and gate-based Variational Quantum Algorithms, offering a self-contained description of the main tools and developments. In both cases, we identified the adoption of hybrid quantum-classical computation as the best strategy for near-term quantum computing.

This review has been written mainly for practitioners in space science and technology with elementary or intermediate knowledge of quantum computing. With this goal in mind, we summarized a large number of research papers, aiming to establish the state-of-the-art and propose new directions for future developments. At the same time, experts

Table 1 This table summarizes the reviewed papers on quantum computing applications in space science and technology. It excludes cited works that focus solely on the theory of quantum computing without direct relevance to space, as well as papers from the space sector that do not involve quantum-related topics. The columns, from left to right, report the articles discussed in Sects. 2, 3, and 4, respectively. Along the rows, papers are organized according to the five major classes of research topics proposed by the Jet Propulsion Lab of NASA.

Synopsis of quantum computing for space applications			
JPL classes	Methodology		
	QA	QML	Q TECH
Earth Science	[75, 76, 78, 86, 89–91, 158–162]§	[77]* [102, 103, 129, 130, 132–135, 140, 144–147, 153, 154, 171–173]	[235, 242, 248]
Astrophysics and Space Science	-	[150, 151, 174]	[250–252, 254]
Planetary Science	[83, 84]	-	[253, 255]
Legacy Technologies	[3, 4, 81, 82, 85, 87, 92–94]	[79, 88]* [152]	[9, 175, 181–183, 208–210, 227–229]
Disruptive Technologies	-	-	[220, 237, 238, 246, 247, 249]

* papers exploiting a gate-based architecture but cited in Sect. 2.

§ papers using QA-based SVMs but cited in Sect. 3.

in quantum computing may benefit from our analysis for identifying practical, real-world problems that extend beyond traditional, idealized test scenarios for novel algorithms and computational proposals.

As anticipated in Sect. 1, here we provide a summary table that may facilitate navigating the review for readers with a background in the space sector. The analyzed studies on quantum applications for space have been discussed or cited in the three main sections of the review, organized by the primary methodologies considered, i.e., quantum annealing in Sect. 2, quantum machine learning in Sect. 3, and integration with other quantum technologies in Sect. 4. Here, we reorganize the screened literature by adopting the classification of the Jet Propulsion Lab of NASA, and classify the reviewed papers into five research fields, as shown in Table 1: (1) Earth sciences, focusing on atmospheric, land, and oceanic studies to predict future changes; (2) planetary sciences, which aim to deepen understanding of planets, satellites, and smaller system bodies; (3) astrophysics and space science, which explore the physics and origins of our galaxy and the universe; (4) legacy technologies, representing technological areas in which the development and upgrade of the instrumentation are crucial for the planning of future missions; and (5) disruptive technologies, which drive system-level innovation in space systems design and implementation.

We provided a pragmatic and realistic assessment of the state-of-the-art, avoiding emphatic and hyped statements on the potential of quantum computing. Novel experimental tests on improved quantum annealing and gate-based hardware will be instrumental in defining the ultimate capability of near-term quantum computing for problems of practical scientific and technological interest, both within and beyond the space sector. Although it is not possible to accurately forecast the development or ultimate technological impact of a complex and rapidly evolving field like quantum computing and its applications to space, we summarize in Table 2 the expected time scales for major anticipated advancements. These range from near-term developments (likely within 3 years), such as the integration of quantum computing with high-performance computing (HPC) and

Table 2 Despite the inherent difficulty in making predictions on emerging technologies, we provide a tentative expected timeline regarding the impact of quantum technologies currently under development. These are classified into three timeframes: *near-term* (likely within 3 years), *mid-term* (likely within 8 years), and *long-term* (likely beyond 8 years). The associated space applications and potential impact provided here are illustrative examples and not aimed to be exhaustive

Estimated timeline for quantum technology deployment in space applications			
Time frame	Quantum technology	Space applications	Impact
Near-term	Integration of HPC and quantum computing	Scheduling and optimization	More efficient assignment of resources
Near-term	Advances in fault-tolerant hardware	RS image classification and accurate simulation of physical processes	Faster and more reliable results
Mid-term	Integration of quantum sensors and quantum computing	Spaceborne data acquisition and multi-messenger astronomy	Improved sensing accuracy and data processing pipeline
Long-term	Distributed quantum computing	Satellite quantum networks	Development of a quantum internet
Long-term	Universal fault-tolerant hardware	Pervasive employment in space applications	New era of space exploration

advances in fault tolerance, to mid-term goals, including the operational integration of quantum computing with quantum sensors, and finally to long-term prospects (likely beyond 8 years), such as distributed quantum computing and universal quantum hardware. To maintain focus, we also provide selected examples of potential space applications and the expected impact of these technological advancements, without aiming to be exhaustive. These timelines reflect expectations shared by parts of the scientific community and governmental organizations, such as the European Flagship program and the Quantum Benchmarking Initiative of DARPA in the USA. However, there is no general consensus, and the development of emerging technologies is always highly non-linear in time, with unexpected breakthroughs driving innovation that are nearly impossible to anticipate.

Our review focused on the scientific side of quantum technologies and quantum computing for space. However, it is worth mentioning the high geopolitical relevance of the field. Recent years have seen a surge in funding directed toward research initiatives that boosted the intersection of quantum technologies and the space sector, reflecting their growing importance in shaping the future of science and innovation. Funding from governments, research institutions, and private ventures is heavily impacting the development of quantum communication networks, quantum sensors, and satellite-based quantum experiments, recognizing their potential to revolutionize data security, navigation, and Earth Observation.

Technological diplomacy and economic competition play a crucial role in the development of quantum computing hardware and software, making joint global blue-sky research alliances unstable and unlikely to form.

Quantum computing and quantum technologies would highly benefit from the long-term vision of advancing global scientific collaboration on the same model as the International Space Station, potentially leading to an international effort to design and deploy a joint quantum science space station [229]. This would represent a key milestone in the integration of quantum technologies into the space infrastructure.

Appendix A: Experimental implementations of quantum computing

Drawing a parallel to one of the most transformative advancements of the last century, we find ourselves in a similar position as in the early days of classical computing. Back then, uncertainty loomed over which materials and technologies would eventually become the state-of-the-art. Ultimately, silicon emerged as the best option, with complementary metal-oxide-semiconductor (CMOS) technology enabling an efficient implementation. This breakthrough set off an avalanche of innovations, revolutionizing daily life and paving the way for current cutting-edge applications such as the Internet of Things, the 6G communication standard, and Artificial Intelligence.

Each quantum computing technology is currently at a different stage of maturity, but all are progressively expanding their computational capabilities. Depending on the number of states involved in each physical quantum unit, the latter is referred to as a qubit (two-level system), qudit (d-level system), or mode (a continuously variable version of qubits).

Here, we mention some notable facts about experimental quantum technologies and introduce the quantities used to benchmark the performance of the quantum processors. The quantum volume [256] is a useful quantity reflecting the largest size of a square circuit that can be successfully run on a specific quantum computer. For instance, a quantum volume of 8 means that 8 qubits are carrying a successful computation of a circuit by involving 8 gate steps. Another relevant parameter is the gate fidelity, especially for two-qubit gates. This parameter is directly related to the maximum circuit depth attainable on a certain technological platform before losing quantum information.

Trapped ions qubits are currently the most successful platform in terms of quantum volume. The qubits are alkali atoms, such as Yb, with prototypes consisting of hundreds of qubits. Trapped ions have a long coherence time and are all-to-all connected, therefore minimizing the depth of the transpiled circuits [257]. Novel 2D arrangements have surpassed the previous 1D geometry. The qubits are operated by lasers in high vacuum and room temperature conditions.

Superconducting qubits have been produced up to over 1000 qubits for gate model architecture and around 5000 qubits for adiabatic quantum computing – with the potential of attaining better performance compared to state-of-the-art classical methods, the so-called "quantum supremacy" [64]. These qubits are based on oscillating LC circuits equipped with a Josephson junction. Such hardware is highly susceptible to physical noise, and cosmic rays are known to produce disturbances to this technology [258]. It requires an mK operating temperature.

In neutral atom platforms [259], qubits are implemented as two electronic states of Rydberg atoms that have been laser-cooled in a vacuum. Quantum error correction has been successfully implemented on this technology [179]. Lasers induce the nK temperature required, so no cryogenics is necessary. Neutral atoms are stabilized in a vacuum and are supported by 3D connectivity featuring a tunable arrangement [260].

Photonic qubits have been commercialized with tens of *qumodes* (continuous variable versions of qubits) potentially suitable for measurement-based quantum computing. Currently available hardware addresses specialized problems, in particular Gaussian boson sampling, but it is planned to be developed for universal quantum computing. The photons carry the quantum information and the circuit can be operated at room temperature.

For what concerns spin qubits in silicon, electron or hole spin encodes the qubit at a temperature of 100-300 mK. Recently, >1 K operability with fidelity of 98-99% was demonstrated. In an optimistic scenario, wafer fabrication makes it reasonable to foresee a scaling to millions of qubits in the long-term future.

Finally, also spin qubits in diamond have been created [21], based on electron states captured by NV-center defects and operated at room temperature.

In gate-model quantum computing, the fragility of quantum states is mitigated by quantum error correction (QEC) methods, which are resource-consuming – a logical qubit may require from 100 to 10,000 physical qubits. When the technology allows to achieve an error rate lower than the error correction rate of the QEC algorithm, then the quantum computer is fault-tolerant. The technologies mentioned above show a different degree of maturity, and it is difficult to predict which of them is going to achieve full development as universal fault-tolerant quantum computers.

Appendix B: A short overview of the theory of classical convolutional neural networks

Computer vision and image classification have been pivotal in driving the current AI revolution, and Convolutional Neural Networks (CNNs) have played a major role. CNNs are a renowned deep learning architecture, tailored to process and classify image or signal data, generally in the form of multidimensional arrays. This architecture is capable of automatically extracting hierarchical and meaningful features from large volumes of image data. CNNs generalize standard feed-forward artificial neural networks and rely on four basic concepts: local connections, shared weights, pooling, and the use of a deep architecture with many layers [124].

Similarly to multi-layer perceptrons, classical CNNs can be trained using backpropagation, and different architecture variations have been proposed to improve their performance, such as LeNet-5 [261], AlexNet [262], and ResNet [263]. As expected, deeper networks tend to achieve higher performance precision, though this comes at the cost of increased optimization complexity and a greater risk of overfitting. All CNN variants share the same building blocks: convolutional layers, pooling layers, and fully-connected layers [264].

In a nutshell, convolutional layers extract feature representations from the input data provided by the previous layer, by using a number of convolutional kernels, each responsible for computing a distinct feature map. Every neuron within these feature maps has a receptive field, that is a localized region of neurons in the previous layer. The feature value at location (i,j) in the k-th feature map of the l-th layer of the network is:

$$z^l_{i,j,k} = \mathbf{w}^{l\,\intercal}_k \mathbf{x}^l_{i,j} + b^l_k\,, \tag{34}$$

where $\mathbf{w}^l_k$ and b^l_k are the weight (kernel) and bias terms, and $\mathbf{x}^l_{i,j}$ is the input patch centered in location (i,j), namely the receptive field. The kernels utilize a weight-sharing mechanism (weights and biases do not depend on the location (i,j)), which simplifies the training process and reduces model complexity.

Then, a non-linear activation function a is applied to enable the detection of non-linear features:

$$a^l_{i,j,k} = a(z^l_{i,j,k})\,. \tag{35}$$

The pooling layers are usually alternated with the convolutional ones and reduce the resolution of the feature maps, contributing to obtaining a shift-invariant feature representation of the input [124]

$$y^{l}_{i,j,k} = \mathrm{pool}(a^{l}_{m,n,k}) \quad \forall (m,n) \in R_{ij}\,, \tag{36}$$

where R_{ij} is a local neighborhood of (i,j). The fully-connected layers are the final layers of the network that perform high-level reasoning. In particular, the last layer is an output layer applied to a high-level feature representation of the input obtained with the previous layers.

CNNs have been devised for and applied to several tasks in machine learning [264]. For our purposes, we consider the standard scenario of supervised learning, e.g. binary or multi-class classification. Here, a CNN is trained by minimizing a loss function such as the cross-entropy loss, typically using variants of stochastic gradient descent [124]. The trainable parameters, collectively denoted as $\boldsymbol{\theta}$, include all the weights and biases of the convolutional and fully-connected layers. The loss function is formally defined as an average over a training set of size N:

$$\mathcal{L} = \frac{1}{N}\sum_{n=1}^{N} \ell(\boldsymbol{\theta}; \mathbf{y}^{(n)}, \mathbf{o}^{(n)}) \tag{37}$$

where $\mathbf{x}^{(n)}$ is the n-th input data, $\mathbf{y}^{(n)}$ the corresponding target label, and $\mathbf{o}^{(n)}$ the predicted label obtained as output of the CNN.

Abbreviations

AEOS, Agile Earth Observation Satellite; AEOSSP, Agile Earth Observation Satellite Scheduling Problem; AQC, Adiabatic Quantum Computation; ASI, Agenzia Spaziale Italiana (in Italian); BP, Barren Plateau; CC, Clique Cover; CNN, Convolutional Neural Network; DQM, Discrete Quadratic Models; DSN, Deep Space Network; EO, Earth Observation; ESA, European Space Agency; FC, Fully Connected; GC, Graph Coloring; GPU, Graphical Processing Unit; HPC, High Performance Computing; HQC-NN, Hybrid Quantum Classical Neural Networks; HSS, Hybrid Solver Service; HUBO, Higher-order Unconstrained Binary Optimization; JSP, Job-shop Scheduling Problem; MBQC, Measurement-based Quantum Computing; ML, Machine Learning; MSE, Mean Square Error; NASA, National Aeronautics and Space Administration; NISQ, Noisy Intermediate-Scale Quantum (Computers); NN, Neural Network; PCA, Principal Component Analysis; PQC, Parametrized Quantum Circuit; PQK, Projected Quantum Kernels; QA, Quantum Annealing; QAOA, Quantum Approximate Optimization Algorithm; QCBM, Quantum Circuit Born Machines; Q-CNN, Quantum Convolutional Neural Network; QML, Quantum Machine Learning; QNN, Quantum Neural Networks; QPU, Quantum Processing Unit; QSVM, Quantum Support Vector Machine; QUBO, Quadratic Unconstrained Binary Optimization; RBM, Restricted Boltzmann Machine; RS, Remote Sensing; SVM, Support Vector Machine; VQA, Variational Quantum Algorithm; VQE, Variational Quantum Eigensolver.

Acknowledgements

The Authors acknowledge the project CQES of the Italian Space Agency (ASI) for having supported this research (grant N. 2023-46-HH.0). The Authors are grateful to Luigi Santamaria Amato of the Italian Space Agency and Alberto Vailati of Università degli Studi di Milano for critically reading the first version of the manuscript and suggesting improvements. The authors thank Francesco Monzani of Università degli Studi di Milano for useful assistance and discussion.

Author contributions

P.T. devised the structure of the review, wrote the main manuscript except Sect. 4, and reviewed the majority of published literature cited in Sects. 1, 2, and 3. R.C. and S.B. helped to screen the literature and review relevant papers for Sects. 2 and 3 and wrote part of the application summaries. A.M. devised and wrote Sect. 4. E.P. supervised the work, wrote parts of Sect. 1, and offered assistance and guidance. All authors reviewed and edited the final manuscript.

Funding information

This research was supported by the project CQES of the Italian Space Agency (ASI), grant N. 2023-46-HH.0.

Data Availability

No datasets were generated or analysed during the current study.

Declarations

Competing interests
The authors declare no competing interests.

Received: 21 February 2025 Accepted: 21 May 2025 Published online: 05 June 2025

References
1. Nielsen MA, Chuang IL. Quantum computation and quantum information. 1st ed. Cambridge: Cambridge University Press; 2000.
2. Preskill J. Quantum computing 40 years later. 2023.
3. Smelyanskiy VN, Rieffel EG, Knysh SI, Williams CP, Johnson MW, Thom MC, Macready WG, Pudenz KL. A near-term quantum computing approach for hard computational problems in space exploration. 2012.
4. Biswas R, Jiang Z, Kechezhi K, Knysh S, Mandrà S, O'Gorman B, Perdomo-Ortiz A, Petukhov A, Realpe-Gómez J, Rieffel E, Venturelli D, Vasko F, Wang Z. A NASA perspective on quantum computing: opportunities and challenges. Parallel Comput. 2017;64:81–98. https://doi.org/10.1016/j.parco.2016.11.002. High-End Computing for Next-Generation Scientific Discovery.
5. Rieffel EG, Hadfield S, Hogg T, Mandrà S, Marshall J, Mossi G, O'Gorman B, Plamadeala E, Tubman NM, Venturelli D, Vinci W, Wang Z, Wilson M, Wudarski F, Biswas R. From ansätze to Z-gates: a NASA view of quantum computing. Amsterdam: IOS Press; 2019. https://doi.org/10.3233/apc190010. http://dx.doi.org/10.3233/APC190010.
6. Noè D, Rocutto L, Moro L, Prati E. Quantum parallel training of a Boltzmann machine on an adiabatic quantum computer. Adv Quantum Technol. 2024;2300330.
7. Kaltenbaek R, Acin A, Bacsardi L, Bianco P, Bouyer P, Diamanti E, Marquardt C, Omar Y, Pruneri V, Rasel E, et al. Quantum technologies in space. Exp Astron. 2021;51(3):1677–94.
8. Belenchia A, Carlesso M, Bayraktar Ö, Dequal D, Derkach I, Gasbarri G, Herr W, Li YL, Rademacher M, Sidhu J, et al. Quantum physics in space. Phys Rep. 2022;951:1–70.
9. Mohageg M, Mazzarella L, Anastopoulos C, Gallicchio J, Hu B-L, Jennewein T, Johnson S, Lin S-Y, Ling A, Marquardt C, et al. The deep space quantum link: prospective fundamental physics experiments using long-baseline quantum optics. EPJ Quantum Technol. 2022;9(1):25.
10. Rieffel EG, Asanjan AA, Alam MS, Anand N, Bernal Neira DE, Block S, Brady LT, Cotton S, Gonzalez Izquierdo Z, Grabbe S, Gustafson E, Hadfield S, Lott PA, Maciejewski FB, Mandrà S, Marshall J, Mossi G, Bauza HM, Saied J, Suri N, Venturelli D, Wang Z, Biswas R. Assessing and advancing the potential of quantum computing: a NASA case study. Future Gener Comput Syst. 2024;160:598–618. https://doi.org/10.1016/j.future.2024.06.012.
11. Alexeev Y, Bacon D, Brown KR, Calderbank R, Carr LD, Chong FT, DeMarco B, Englund D, Farhi E, Fefferman B, et al. Quantum computer systems for scientific discovery. PRX Quantum. 2021;2(1):017001.
12. Albash T, Lidar DA. Adiabatic quantum computation. Rev Mod Phys. 2018;90:015002. https://doi.org/10.1103/RevModPhys.90.015002.
13. Briegel HJ, Browne DE, Dür W, Raussendorf R, Nest M. Measurement-based quantum computation. Nat Phys. 2009;5(1):19–26.
14. Ladd TD, Jelezko F, Laflamme R, Nakamura Y, Monroe C, O'Brien JL. Quantum computers. Nature. 2010;464(7285):45–53.
15. Huang H-L, Wu D, Fan D, Zhu X. Superconducting quantum computing: a review. Sci China Inf Sci. 2020;63:1–32.
16. Manovitz T, Shapira Y, Gazit L, Akerman N, Ozeri R. Trapped-ion quantum computer with robust entangling gates and quantum coherent feedback. PRX Quantum. 2022;3(1):010347.
17. Cong I, Levine H, Keesling A, Bluvstein D, Wang S-T, Lukin MD. Hardware-efficient, fault-tolerant quantum computation with Rydberg atoms. Phys Rev X. 2022;12(2):021049.
18. Rotta D, Sebastiano F, Charbon E, Prati E. Quantum information density scaling and qubit operation time constraints of CMOS silicon-based quantum computer architectures. npj Quantum Inf. 2017;3(1):26.
19. Ferraro E, Prati E. Is all-electrical silicon quantum computing feasible in the long term? Phys Lett A. 2020;384(17):126352.
20. De Michielis M, Ferraro E, Prati E, Hutin L, Bertrand B, Charbon E, Ibberson DJ, Gonzalez-Zalba MF. Silicon spin qubits from laboratory to industry. J Phys D, Appl Phys. 2023;56(36):363001.
21. Pezzagna S, Meijer J. Quantum computer based on color centers in diamond. Appl Phys Rev. 2021;8(1).
22. Flamini F, Spagnolo N, Sciarrino F. Photonic quantum information processing: a review. Rep Prog Phys. 2018;82(1):016001.
23. Preskill J. Quantum computing in the NISQ era and beyond. Quantum. 2018;2:79. https://doi.org/10.22331/q-2018-08-06-79.
24. Apolloni B, Cesa-Bianchi N, De Falco D. A numerical implementation of "quantum annealing". In: Stochastic processes, physics and geometry: proceedings of the ascona-locarno conference. 1990. p. 97–111.
25. Yarkoni S, Raponi E, Bäck T, Schmitt S. Quantum annealing for industry applications: introduction and review. Rep Prog Phys. 2022;85(10):104001. https://doi.org/10.1088/1361-6633/ac8c54.
26. Rønnow TF, Wang Z, Job J, Boixo S, Isakov SV, Wecker D, Martinis JM, Lidar DA, Troyer M. Defining and detecting quantum speedup. Science. 2014;345(6195):420–4. https://doi.org/10.1126/science.1252319.
27. Auffèves A. Quantum technologies need a quantum energy initiative. PRX Quantum. 2022;3:020101. https://doi.org/10.1103/PRXQuantum.3.020101.
28. Kim Y, Eddins A, Anand S, Wei KX, Berg E, Rosenblatt S, Nayfeh H, Wu Y, Zaletel M, Temme K, Kandala A. Evidence for the utility of quantum computing before fault tolerance. Nature. 2023;618:500–5. https://doi.org/10.1038/s41586-023-05955-1.
29. Zufferey N, Amstutz P, Giaccari P. Graph colouring approaches for a satellite range scheduling problem. J Sched. 2008;11(4):263–77. https://doi.org/10.1007/s10951-008-0066-8.
30. Marx D. Graph colouring problems and their applications in scheduling. Period Polytech, Electr Eng. 2004;48(1–2):11–6.

31. Venturelli D, Marchand DJJ, Rojo G. Quantum annealing implementation of job-shop scheduling. 2016.
32. Glover F, Kochenberger G, Du Y. A tutorial on formulating and using QUBO models. 2019.
33. Lucas A. Ising formulations of many NP problems. Front Phys. 2014;2:5.
34. Farhi E, Goldstone J, Gutmann S. A quantum approximate optimization algorithm. 2014.
35. Hadfield S, Wang Z, O'Gorman B, Rieffel E, Venturelli D, Biswas R. From the quantum approximate optimization algorithm to a quantum alternating operator ansatz. Algorithms. 2019;12(2):34. https://doi.org/10.3390/a12020034.
36. Blekos K, Brand D, Ceschini A, Chou C-H, Li R-H, Pandya K, Summer A. A review on quantum approximate optimization algorithm and its variants. 2023.
37. Zhou L, Wang S-T, Choi S, Pichler H, Lukin MD. Quantum approximate optimization algorithm: performance, mechanism, and implementation on near-term devices. Phys Rev X. 2020;10:021067. https://doi.org/10.1103/PhysRevX.10.021067.
38. Corli S, Dragoni D, Proietti M, Dispenza M, Cavazzoni C, Prati E. A max K-cut implementation for QAOA in the measurement based quantum computing formalism. In: 2023 IEEE international conference on Quantum Computing and Engineering (QCE). vol. 2. Los Alamitos: IEEE; 2023. p. 284–5.
39. Pecci G, Wang R, Torta P, Mbeng GB, Santoro G. Beyond quantum annealing: optimal control solutions to maxcut problems. Quantum Sci Technol. 2024;9(4):045013. https://doi.org/10.1088/2058-9565/ad60f2.
40. Zentilini P, Corli S, Prati E. Emulating QAOA via graph neural networks. In: 2024 IEEE international conference on Quantum Computing and Engineering (QCE). vol. 2. Los Alamitos: IEEE; 2024. p. 472–3.
41. Cerezo M, Arrasmith A, Babbush R, Benjamin S, Endo S, Fujii K, McClean JR, Mitarai K, Yuan X, Cincio L, Coles P. Variational Quantum Algorithms. Nat Rev Phys. 2021.
42. Bharti K, Cervera-Lierta A, Kyaw TH, Haug T, Alperin-Lea S, Anand A, Degroote M, Heimonen H, Kottmann JS, Menke T, Mok W-K, Sim S, Kwek L-C, Aspuru-Guzik A. Noisy intermediate-scale quantum algorithms. Rev Mod Phys. 2022;94:015004. https://doi.org/10.1103/RevModPhys.94.015004.
43. Finnila AB, Gomez MA, Sebenik C, Stenson C, Doll JD. Quantum annealing: a new method for minimizing multidimensional functions. Chem Phys Lett. 1994;219:343.
44. Kadowaki T, Nishimori H. Quantum annealing in the transverse Ising model. Phys Rev E. 1998;58:5355–63. https://doi.org/10.1103/PhysRevE.58.5355.
45. Santoro GE, Martoňák R, Tosatti E, Car R. Theory of quantum annealing of an Ising spin glass. Science. 2002;295:2427.
46. Santoro GE, Tosatti E. Optimization using quantum mechanics: quantum annealing through adiabatic evolution. J Phys A, Math Gen. 2006;39:393–431. https://doi.org/10.1088/0305-4470/39/36/R01.
47. Farhi E, Goldstone J, Gutmann S, Lapan J, Lundgren A, Preda D. A quantum adiabatic evolution algorithm applied to random instances of an NP-complete problem. Science. 2001;292(5516):472–5. https://doi.org/10.1126/science.1057726.
48. Morita S, Nishimori H. Mathematical foundation of quantum annealing. J Math Phys. 2008;**49**(12).
49. Aharonov D, Dam W, Kempe J, Landau Z, Lloyd S, Regev O. Adiabatic quantum computation is equivalent to standard quantum computation. In: 45th annual IEEE symposium on foundations of computer science. 2004. p. 42–51. https://doi.org/10.1109/FOCS.2004.8.
50. Vinci W, Lidar DA. Non-stoquastic Hamiltonians in quantum annealing via geometric phases. npj Quantum Inf. 2017;**3**(1). https://doi.org/10.1038/s41534-017-0037-z.
51. Scholl P, Schuler M, Williams HJ, et al. Quantum simulation of 2d antiferromagnets with hundreds of Rydberg atoms. Nature. 2021;595:233–8. https://doi.org/10.1038/s41586-021-03585-1.
52. Reichardt BW, Paetznick A, Aasen D, Basov I, Bello-Rivas JM, Bonderson P, Chao R, Dam W, Hastings MB, Paz A, et al. Logical computation demonstrated with a neutral atom quantum processor. 2024. arXiv preprint arXiv:2411.11822.
53. Ebadi S, Wang TT, Levine H, et al. Quantum phases of matter on a 256-atom programmable quantum simulator. Nature. 2021;595:227–32. https://doi.org/10.1038/s41586-021-03582-4.
54. D-Wave Systems Inc. Advantage processor overview. Technical Report 14-1058A, D-Wave Systems. 2020. https://www.dwavesys.com/media/3xvdipcn/14-1058a-a_advantage_processor_overview.pdf. Accessed: yyyy-mm-dd.
55. D-Wave Systems Inc. Annealing Implementation and Controls. https://docs.dwavesys.com/docs/latest/c_qpu_annealing.html#id1. Accessed yyyy-mm-dd.
56. Boothby K, Bunyk P, Raymond J, Roy A. Next-generation topology of D-wave quantum processors. 2020.
57. Cai J, Macready WG, Roy A. A practical heuristic for finding graph minors. 2014. https://arxiv.org/abs/1406.2741.
58. D-Wave Systems Inc. D-wave hybrid solver service: an overview. Technical Report 14-1039A, D-Wave Systems. 2020. https://www.dwavesys.com/media/4bnpi53x/14-1039a-b_d-wave_hybrid_solver_service_an_overview.pdf.
59. D-Wave Systems Inc. Hybrid solver for constrained quadratic models. Technical Report 14-1055A, D-Wave Systems. 2021. https://www.dwavesys.com/media/rldh2ghw/14-1055a-a_hybrid_solver_for_constrained_quadratic_models.pdf.
60. Osaba E, Miranda-Rodriguez P. D-wave's nonlinear-program hybrid solver: description and performance analysis. 2024. https://arxiv.org/abs/2410.07980.
61. Higham CF, Bedford A. Quantum deep learning by sampling neural nets with a quantum annealer. Sci Rep. 2023;13:3939. https://doi.org/10.1038/s41598-023-30910-7.
62. Rønnow TF, Wang Z, Job J, Boixo S, Isakov SV, Wecker D, Martinis JM, Lidar DA, Troyer M. Defining and detecting quantum speedup. Science. 2014;345(6195):420–4.
63. Moro L, Prati E. Anomaly detection speed-up by quantum restricted Boltzmann machines. Commun Phys. 2023;6(1):269.
64. King AD, Nocera A, Rams MM, Dziarmaga J, Wiersema R, Bernoudy W, Raymond J, Kaushal N, Heinsdorf N, Harris R, Boothby K, Altomare F, Berkley AJ, Boschnak M, Chern K, Christiani H, Cibere S, Connor J, Dehn MH, Deshpande R, Ejtemaee S, Farré P, Hamer K, Hoskinson E, Huang S, Johnson MW, Kortas S, Ladizinsky E, Lai T, Lanting T, Li R, MacDonald AJR, Marsden G, McGeoch CC, Molavi R, Neufeld R, Norouzpour M, Oh T, Pasvolsky J, Poitras P, Poulin-Lamarre G, Prescott T, Reis M, Rich C, Samani M, Sheldan B, Smirnov A, Sterpka E, Clavera BT, Tsai N, Volkmann M, Whiticar A, Whittaker JD, Wilkinson W, Yao J, Yi TJ, Sandvik AW, Alvarez G, Melko RG, Carrasquilla J, Franz M, Amin MH. 2024. Computational supremacy in quantum simulation. https://arxiv.org/abs/2403.00910.
65. Chen H, Lidar DA. Why and when pausing is beneficial in quantum annealing. Phys Rev Appl. 2020;**14**(1). https://doi.org/10.1103/physrevapplied.14.014100.

66. D-Wave Systems Inc. Reverse Quantum Annealing for Local Refinement of Solutions. https://www.dwavesys.com/media/5hsjmvom/14-1018a-a_reverse_quantum_annealing_for_local_refinement_of_solutions.pdf.
67. Dauzère-Pérès S, Ding J, Shen L, Tamssaouet K. The flexible job shop scheduling problem: a review. Eur J Oper Res. 2024;314(2):409–32. https://doi.org/10.1016/j.ejor.2023.05.017.
68. Jensen TR, Toft B. Graph coloring problems. New York: Wiley; 2011.
69. Gromicho JAS, van Hoorn JJ, Saldanha-da-Gama F, Timmer GT. Solving the job-shop scheduling problem optimally by dynamic programming. Comput Oper Res. 2012;39(12):2968–77. https://doi.org/10.1016/j.cor.2012.02.024.
70. Carugno C, Ferrari Dacrema M, Cremonesi P. Evaluating the job shop scheduling problem on a D-wave quantum annealer. Sci Rep. 2022;12(1):6539. https://doi.org/10.1038/s41598-022-10169-0.
71. Kurowski K, Węglarz J, Subocz M, Różycki R, Waligóra G. Hybrid quantum annealing heuristic method for solving job shop scheduling problem. In: Krzhizhanovskaya VV, Závodszky G, Lees MH, Dongarra JJ, Sloot PMA, Brissos S, Teixeira J, editors. Computational science – ICCS 2020. Cham: Springer; 2020. p. 502–15.
72. Torta P, Leone L, Casati R, Prati E. Neural quantum annealing for real-world quadratic unconstrained binary optimization. In: 2024 IEEE international conference on Quantum Computing and Engineering (QCE). vol. 02. 2024. p. 514–5. https://doi.org/10.1109/QCE60285.2024.10382.
73. Błażewicz J, Pesch E, Sterna M. The disjunctive graph machine representation of the job shop scheduling problem. Eur J Oper Res. 2000;127(2):317–31. https://doi.org/10.1016/S0377-2217(99)00486-5.
74. Wang X, Wu G, Xing L, Pedrycz W. Agile Earth observation satellite scheduling over 20 years: formulations, methods, and future directions. IEEE Syst J. 2021;15(3):3881–92. https://doi.org/10.1109/jsyst.2020.2997050.
75. Stollenwerk T, Michaud V, Lobe E, Picard M, Basermann A, Botter T. Agile Earth observation satellite scheduling with a quantum annealer. IEEE Trans Aerosp Electron Syst. 2021;57(5):3520–8. https://doi.org/10.1109/TAES.2021.3088490.
76. Makarov A, Pérez-Herradón C, Franceschetto G, Taddei MM, Osaba E, Barrio Cabello P, Villar-Rodriguez E, Oregi I. Quantum optimization methods for satellite mission planning. IEEE Access. 2024;12:71808–20. https://doi.org/10.1109/ACCESS.2024.3402990.
77. Rainjonneau S, Tokarev I, Iudin S, Rayaprolu S, Pinto K, Lemtiuzhnikova D, Koblan M, Barashov E, Kordzanganeh M, Pflitsch M, Melnikov A. Quantum algorithms applied to satellite mission planning for Earth observation. IEEE J Sel Top Appl Earth Obs Remote Sens. 2023;16:7062–75. https://doi.org/10.1109/jstars.2023.3287154.
78. Delilbasic A, Le Saux B, Riedel M, Michielsen K, Cavallaro G. Reverse quantum annealing for hybrid quantum-classical satellite mission planning. In: IGARSS 2024 - 2024 IEEE international geoscience and remote sensing symposium. 2024. p. 432–6. https://doi.org/10.1109/IGARSS53475.2024.10640974.
79. Quetschlich N, Koch V, Burgholzer L, Wille R. A hybrid classical quantum computing approach to the satellite mission planning problem. In: 2023 IEEE international conference on Quantum Computing and Engineering (QCE). vol. 1. Los Alamitos: IEEE; 2023. p. 642–7.
80. Clement BJ, Johnston MD. The deep space network scheduling problem. In: Proceedings of the 17th Conference on Innovative Applications of Artificial Intelligence (IAAI'05). Pittsburgh: AAAI Press; 2005. p. 1514–20. https://fileadmin.cs.lth.se/ai/Proceedings/aaai05/PAPERS/IAAI05-009.pdf.
81. Guillaume A, Goh EY, Johnston MD, Wilson BD, Ramanan A, Tibble F, Lackey B. Deep space network scheduling using quantum annealing. IEEE Trans Quantum Eng. 2022;3:1–13.
82. Wilson B, Goh E, Guillaume A, Alimo R, Claudet T, Venkataram H. Automating antenna scheduling problems using quantum computing and deep reinforcement learning. In: IGARSS 2022 - 2022 IEEE international geoscience and remote sensing symposium. 2022. p. 4915–8. https://doi.org/10.1109/IGARSS46834.2022.9884342.
83. Tran TT, Wang Z, Do M, Rieffel EG, Frank J, O'Gorman B, Venturelli D, Beck JC. Explorations of quantum-classical approaches to scheduling a Mars lander activity problem. In: Workshops at the thirtieth AAAI conference on artificial intelligence. 2016.
84. Tran TT, Do MB, Rieffel EG, Frank JD, Wang Z, O'Gorman B, Venturelli D, Beck JC. A hybrid quantum-classical approach to solving scheduling problems. In: Symposium on combinatorial search. 2016. https://api.semanticscholar.org/CorpusID:9725684.
85. Perdomo-Ortiz A, Fluegemann J, Narasimhan S, Biswas R, Smelyanskiy VN. A quantum annealing approach for fault detection and diagnosis of graph-based systems. Eur Phys J Spec Top. 2015;224:131–48.
86. Rocutto L, Maronese M, Dragoni D, Cavalli A, Cavazzoni C. Comparing adiabatic quantum computers for satellite images feature extraction. Future Gener Comput Syst. 2024;159:105–13.
87. Kim M, Venturelli D, Jamieson K. Leveraging quantum annealing for large mimo processing in centralized radio access networks. In: Proceedings of the ACM special interest group on data communication. 2019. p. 241–55.
88. Al-Hraishawi H, Rehman J, Chatzinotas S. Quantum optimization algorithm for Leo satellite communications based on cell-free massive mimo. In: 2023 IEEE International Conference on Communications Workshops (ICC Workshops). Los Alamitos: IEEE; 2023. p. 1759–64.
89. Otgonbaatar S, Datcu M. A quantum annealer for subset feature selection and the classification of hyperspectral images. IEEE J Sel Top Appl Earth Obs Remote Sens. 2021;14:7057–65. https://doi.org/10.1109/JSTARS.2021.3095377.
90. Bass G, Tomlin C, Kumar V, Rihaczek P, Dulny J. Heterogeneous quantum computing for satellite constellation optimization: solving the weighted k-clique problem. Quantum Sci Technol. 2018;3(2):024010. https://doi.org/10.1088/2058-9565/aaadc2.
91. Kim J, Kwak Y, Jung S, Kim J-H. Quantum scheduling for millimeter-wave observation satellite constellation. In: 2021 IEEE VTS 17th Asia Pacific Wireless Communications Symposium (APWCS). Los Alamitos: IEEE; 2021. p. 1–5.
92. Dinh TQ, Dau SH, Lagunas E, Chatzinotas S, Nguyen DN, Hoang DT. Quantum annealing for complex optimization in satellite communication systems. IEEE Internet Things J. 2024.
93. Stollenwerk T, O'Gorman B, Venturelli D, Mandra S, Rodionova O, Ng H, Sridhar B, Rieffel EG, Biswas R. Quantum annealing applied to de-conflicting optimal trajectories for air traffic management. IEEE Trans Intell Transp Syst. 2019;21(1):285–97.
94. Carbone A, De Grossi F, Spiller D. Cutting-edge trajectory optimization through quantum annealing. Appl Sci. 2023;13(23):12853.
95. Yang Y, Newsam S. Bag-of-visual-words and spatial extensions for land-use classification. In: Proceedings of the 18th SIGSPATIAL international conference on advances in geographic information systems. GIS '10. New York: Assoc. Comput. Mach.; 2010. p. 270–9. https://doi.org/10.1145/1869790.1869829.

96. Cheng G, Han J, Lu X. Remote sensing image scene classification: benchmark and state of the art. Proc IEEE. 2017;105(10):1865–83. https://doi.org/10.1109/JPROC.2017.2675998.
97. Kansakar P, Hossain F. A review of applications of satellite Earth observation data for global societal benefit and stewardship of planet Earth. Space Policy. 2016;36:46–54.
98. Zhang L, Zhang L, Du B. Deep learning for remote sensing data: a technical tutorial on the state of the art. IEEE Geosci Remote Sens Mag. 2016;4:22–40.
99. Zhu XX, Tuia D, Mou L, Xia G-S, Zhang L, Xu F, Fraundorfer F. Deep learning in remote sensing: a comprehensive review and list of resources. IEEE Geosci Remote Sens Mag. 2017;5(4):8–36. https://doi.org/10.1109/MGRS.2017.2762307.
100. Cheng G, Xie X, Han J, Guo L, Xia G-S. Remote sensing image scene classification meets deep learning: challenges, methods, benchmarks, and opportunities. IEEE J Sel Top Appl Earth Obs Remote Sens. 2020;13:3735–56. https://doi.org/10.1109/JSTARS.2020.3005403.
101. Reichstein M, Camps-Valls G, Stevens B, Jung M, Denzler J, Carvalhais N, Prabhat. Deep learning and process understanding for data-driven Earth system science. Nature. 2019;566:195–204. https://doi.org/10.1038/s41586-019-0912-1.
102. Otgonbaatar S, Kranzlmüller D. Exploiting the quantum advantage for satellite image processing: review and assessment. IEEE Trans Quantum Eng. 2023:1–9. https://doi.org/10.1109/TQE.2023.3338970.
103. Nammouchi A, Kassler A, Theorachis A. Quantum machine learning in climate change and sustainability: a review. 2023. https://arxiv.org/abs/2310.09162.
104. Torta P. Quantum approximate optimization algorithm and variational quantum computing: from binary neural networks to ground state preparation. PhD thesis. Scuola Internazionale Superiore di Studi Avanzati (SISSA), Trieste, Italy; (April 2024). Supervisor: giuseppe Ernesto Santoro. https://iris.sissa.it/handle/20.500.11767/138170.
105. Benedetti M, Lloyd E, Sack S, Fiorentini M. Parameterized quantum circuits as machine learning models. Quantum Sci Technol. 2019;4(4):043001. https://doi.org/10.1088/2058-9565/ab4eb5.
106. Sørensen A, Mølmer K. Entanglement and quantum computation with ions in thermal motion. Phys Rev A. 2000;**62**(2). https://doi.org/10.1103/physreva.62.022311.
107. Peruzzo A, McClean J, Shadbolt P, Yung M-H, Zhou X-Q, Love PJ, Aspuru-Guzik A, O'Brien JL. A variational eigenvalue solver on a photonic quantum processor. Nat Commun. 2014;5(1):4213. https://doi.org/10.1038/ncomms5213.
108. Agliardi G, Prati E. Quantum data encoding as a distinct abstraction layer in the design of quantum circuits. Quantum Sci Technol. 2025;10(2):025008. https://doi.org/10.1088/2058-9565/ada6f8.
109. Schuld M, Bocharov A, Svore KM, Wiebe N. Circuit-centric quantum classifiers. Phys Rev A. 2020;**101**(3). https://doi.org/10.1103/physreva.101.032308.
110. Schuld M, Sweke R, Meyer JJ. Effect of data encoding on the expressive power of variational quantum-machine-learning models. Phys Rev A. 2021;103:032430. https://doi.org/10.1103/PhysRevA.103.032430.
111. Aleksandrowicz G, Alexander T, Barkoutsos P, Bello L, Ben-Haim Y, Bucher D, Cabrera-Hernández FJ, Carballo-Franquis J, Chen A, Chen C-F, et al. Qiskit: an open-source framework for quantum computing. 2019. https://doi.org/10.5281/zenodo.2562111.
112. Bergholm V, Izaac J, Schuld M, Gogolin C, Alam MS, Ahmed S, Arrazola JM, Blank C, Delgado A, Jahangiri S, McKiernan K, Meyer JJ, Niu Z, Száva A, Killoran N. PennyLane: automatic differentiation of hybrid quantum-classical computations. 2020.
113. Wierichs D, Izaac J, Wang C, Lin CY-Y. General parameter-shift rules for quantum gradients. Quantum. 2022;6:677. https://doi.org/10.22331/q-2022-03-30-677.
114. Cai Z, Babbush R, Benjamin SC, Endo S, Huggins WJ, Li Y, McClean JR, O'Brien TE. Quantum error mitigation. Rev Mod Phys. 2023;**95**(4). https://doi.org/10.1103/revmodphys.95.045005.
115. Anschuetz ER, Kiani BT. Quantum variational algorithms are swamped with traps. Nat Commun. 2022;**13**(1). https://doi.org/10.1038/s41467-022-35364-5.
116. Bittel L, Kliesch M. Training variational quantum algorithms is NP-hard. Phys Rev Lett. 2021;127:120502. https://doi.org/10.1103/PhysRevLett.127.120502.
117. Arrasmith A, Holmes Z, Cerezo M, Coles PJ. Equivalence of quantum barren plateaus to cost concentration and narrow gorges. Quantum Sci Technol. 2022;7(4):045015. https://doi.org/10.1088/2058-9565/ac7d06.
118. Cerezo M, Larocca M, García-Martín D, Diaz NL, Braccia P, Fontana E, Rudolph MS, Bermejo P, Ijaz A, Thanasilp S, Anschuetz ER, Holmes Z. Does provable absence of barren plateaus imply classical simulability? Or, why we need to rethink variational quantum computing. 2024. https://arxiv.org/abs/2312.09121.
119. Ragone M, Bakalov BN, Sauvage F, Kemper AF, Ortiz Marrero C, Larocca M, Cerezo M. A Lie algebraic theory of barren plateaus for deep parameterized quantum circuits. Nat Commun. 2024;**15**(1). https://doi.org/10.1038/s41467-024-49909-3.
120. Fontana E, Herman D, Chakrabarti S, Kumar N, Yalovetzky R, Heredge J, Sureshbabu SH, Pistoia M. Characterizing barren plateaus in quantum ansätze with the adjoint representation. Nat Commun. 2024;**15**(1). https://doi.org/10.1038/s41467-024-49910-w.
121. Nest MV. Classical simulation of quantum computation, the Gottesman-Knill theorem, and slightly beyond. 2009.
122. Collura M, Lami G, Ranabhat N, Santini A. In: Tensor network techniques for quantum computation. SISSA medialab S.r.l. Trieste, Italy. 2024. https://doi.org/10.22323/9788898587049. https://library.oapen.org/handle/20.500.12657/96026.
123. Carleo G, Troyer M. Solving the quantum many-body problem with artificial neural networks. Science. 2017;355(6325):602–6. https://doi.org/10.1126/science.aag2302.
124. LeCun Y, Bengio Y, Hinton G. Deep learning. Nature. 2015;521(7553):436–44.
125. Cong I, Choi S, Lukin MD. Quantum convolutional neural networks. Nat Phys. 2019;15(12):1273–8. https://doi.org/10.1038/s41567-019-0648-8.
126. Henderson M, Shakya S, Pradhan S, Cook T. Quanvolutional neural networks: powering image recognition with quantum circuits. Quantum Mach Intell. 2020;2(2):1–9. https://doi.org/10.1007/s42484-020-00012-y.
127. Hur T, Kim L, Park DK. Quantum convolutional neural network for classical data classification. Quantum Mach Intell. 2022;4(1):3.

128. Monaco S, Kiss O, Mandarino A, Vallecorsa S, Grossi M. Quantum phase detection generalization from marginal quantum neural network models. Phys Rev B. 2023;107:081105. https://doi.org/10.1103/PhysRevB.107.L081105.
129. Henderson M, Gallina J, Brett M. Methods for accelerating geospatial data processing using quantum computers. Quantum Mach Intell. 2021;3(1):4.
130. Chang SY, Le Saux B, Vallecorsa S, Grossi M. Quantum convolutional circuits for Earth observation image classification. In: IGARSS 2022-2022 IEEE international geoscience and remote sensing symposium. Los Alamitos: IEEE; 2022. p. 4907–10.
131. Vidal G. Class of quantum many-body states that can be efficiently simulated. Phys Rev Lett. 2008;101(11):110501.
132. Zaidenberg DA, Sebastianelli A, Spiller D, Le Saux B, Ullo SL. Advantages and bottlenecks of quantum machine learning for remote sensing. In: 2021 IEEE International Geoscience and Remote Sensing Symposium IGARSS. 2021. p. 5680–3. https://doi.org/10.1109/IGARSS47720.2021.9553133.
133. Otgonbaatar S, Datcu M. Classification of remote sensing images with parameterized quantum gates. IEEE Geosci Remote Sens Lett. 2021;19:1–5.
134. Gawron P, Lewiński S. Multi-spectral image classification with quantum neural network. In: IGARSS 2020-2020 IEEE international geoscience and remote sensing symposium. Los Alamitos: IEEE; 2020. p. 3513–6.
135. Sebastianelli A, Zaidenberg DA, Spiller D, Saux BL, Ullo SL. On circuit-based hybrid quantum neural networks for remote sensing imagery classification. IEEE J Sel Top Appl Earth Obs Remote Sens. 2022;15:565–80. https://doi.org/10.1109/JSTARS.2021.3134785.
136. Helber P, Bischke B, Dengel A, Borth D. Eurosat: a novel dataset and deep learning benchmark for land use and land cover classification. IEEE J Sel Top Appl Earth Obs Remote Sens. 2019;12(7):2217–26. https://doi.org/10.1109/JSTARS.2019.2918242.
137. Kandala A, Mezzacapo A, Temme K, Takita M, Brink M, Chow JM, Gambetta JM. Hardware-efficient variational quantum eigensolver for small molecules and quantum magnets. Nature. 2017;549:242–6.
138. Meyer JJ, Mularski M, Gil-Fuster E, Mele AA, Arzani F, Wilms A, Eisert J. Exploiting symmetry in variational quantum machine learning. PRX Quantum. 2023;4:010328. https://doi.org/10.1103/PRXQuantum.4.010328.
139. Pérez-Salinas A, Cervera-Lierta A, Gil-Fuster E, Latorre JI. Data re-uploading for a universal quantum classifier. Quantum. 2020;4:226. https://doi.org/10.22331/q-2020-02-06-226.
140. Zollner JM, Walther P, Werner M. Satellite image representations for quantum classifiers. Datenbank Spektrum. 2024;24:33–41. https://doi.org/10.1007/s13222-024-00464-7.
141. Farhi E, Neven H. Classification with quantum neural networks on near term processors. 2018. arXiv preprint arXiv:1802.06002.
142. Grant E, Benedetti M, Cao S, Hallam A, Lockhart J, Stojevic V, Green AG, Severini S. Hierarchical quantum classifiers. npj Quantum Inf. 2018;4(1):65.
143. Lazzarin M, Galli DE, Prati E. Multi-class quantum classifiers with tensor network circuits for quantum phase recognition. Phys Lett A. 2022;434:128056.
144. Otgonbaatar S, Schwarz G, Datcu M, Kranzlmüller D. Quantum transfer learning for real-world, small, and high-dimensional remotely sensed datasets. IEEE J Sel Top Appl Earth Obs Remote Sens. 2023;16:9223–30. https://doi.org/10.1109/JSTARS.2023.3316306.
145. Papa L, Sebastianelli A, Meoni G, Amerini I. On the impact of key design aspects in simulated hybrid quantum neural networks for Earth observation. 2024. arXiv preprint arXiv:2410.08677.
146. Fan F, Shi Y, Zhu XX. Land cover classification from sentinel-2 images with quantum-classical convolutional neural networks. IEEE J Sel Top Appl Earth Obs Remote Sens. 2024;17:12477–89. https://doi.org/10.1109/JSTARS.2024.3411670.
147. Fan F, Shi Y, Zhu XX. Earth observation data classification with quantum-classical convolutional neural network. In: IGARSS 2022 - 2022 IEEE international geoscience and remote sensing symposium. 2022. p. 191–4. https://doi.org/10.1109/IGARSS46834.2022.9883949.
148. Le PQ, Dong F, Hirota K. A flexible representation of quantum images for polynomial preparation, image compression, and processing operations. Quantum Inf Process. 2011;10:63–84. https://doi.org/10.1007/s11128-010-0177-y.
149. Amankwah MG, Camps D, Bethel EW, et al. Quantum pixel representations and compression for n-dimensional images. Sci Rep. 2022;12:7712. https://doi.org/10.1038/s41598-022-11024-y.
150. Alomari A, Kumar SAP. Hybrid classical-quantum neural network for improving space weather detection and early warning alerts. In: 2023 IEEE Cognitive Communications for Aerospace Applications Workshop (CCAAW). 2023. p. 1–6. https://doi.org/10.1109/CCAAW57883.2023.10219316.
151. Sushmit MM, Mahbubul IM. Forecasting solar irradiance with hybrid classical–quantum models: a comprehensive evaluation of deep learning and quantum-enhanced techniques. Energy Convers Manag. 2023;294:117555.
152. Fadli S, Rawal BS. Hybrid quantum-classical machine learning for near real-time space to ground communication of ISS lightning imaging sensor data. In: 2023 IEEE 13th annual computing and Communication Workshop and Conference (CCWC). 2023. p. 0114–22. https://doi.org/10.1109/CCWC57344.2023.10099338.
153. Otgonbaatar S, Datcu M. Natural embedding of the Stokes parameters of polarimetric synthetic aperture radar images in a gate-based quantum computer. IEEE Trans Geosci Remote Sens. 2022;60:1–8. https://doi.org/10.1109/TGRS.2021.3110056.
154. Rodriguez-Grasa P, Farzan-Rodriguez R, Novelli G, Ban Y, Sanz M. Satellite image classification with neural quantum kernels. Mach Learn: Sci Technol. 2025;6(1):015043. https://doi.org/10.1088/2632-2153/ada86c.
155. Mountrakis G, Im J, Ogole C. Support vector machines in remote sensing: a review. ISPRS J Photogramm Remote Sens. 2011;66(3):247–59. https://doi.org/10.1016/j.isprsjprs.2010.11.001.
156. Burges CJC. A tutorial on support vector machines for pattern recognition. Data Min Knowl Discov. 1998;2(2):121–67. https://doi.org/10.1023/A:1009715923555.
157. Willsch D, Willsch M, De Raedt H, Michielsen K. Support vector machines on the d-wave quantum annealer. Comput Phys Commun. 2020;248:107006. https://doi.org/10.1016/j.cpc.2019.107006.
158. Cavallaro G, Willsch D, Willsch M, Michielsen K, Riedel M. Approaching remote sensing image classification with ensembles of support vector machines on the d-wave quantum annealer. In: IGARSS 2020-2020 IEEE international geoscience and remote sensing symposium. Los Alamitos: IEEE; 2020. p. 1973–6.

159. Delilbasic A, Cavallaro G, Willsch M, Melgani F, Riedel M, Michielsen K. Quantum support vector machine algorithms for remote sensing data classification. In: 2021 IEEE International Geoscience and Remote Sensing Symposium IGARSS. Los Alamitos: IEEE; 2021. p. 2608–11.
160. Otgonbaatar S, Datcu M. Assembly of a coreset of Earth observation images on a small quantum computer. Electronics. 2021;10(20):2482.
161. Pasetto E, Delilbasic A, Cavallaro G, Willsch M, Melgani F, Riedel M, Michielsen K. Quantum support vector regression for biophysical variable estimation in remote sensing. In: IGARSS 2022 - 2022 IEEE international geoscience and remote sensing symposium. 2022. p. 4903–6. https://doi.org/10.1109/IGARSS46834.2022.9883963.
162. Delilbasic A, Saux B, Riedel M, Michielsen K, Cavallaro G. A single-step multiclass SVM based on quantum annealing for remote sensing data classification. IEEE J Sel Top Appl Earth Obs Remote Sens 2023:1–12. https://doi.org/10.1109/JSTARS.2023.3336926.
163. Rebentrost P, Mohseni M, Lloyd S. Quantum support vector machine for big data classification. Phys Rev Lett. 2014;113:130503. https://doi.org/10.1103/PhysRevLett.113.130503.
164. Harrow AW, Hassidim A, Lloyd S. Quantum algorithm for linear systems of equations. Phys Rev Lett. 2009;103(15):150502.
165. Schuld M, Killoran N. Quantum machine learning in feature Hilbert spaces. Phys Rev Lett. 2019;**122**(4). https://doi.org/10.1103/physrevlett.122.040504.
166. Havlíček V, Córcoles AD, Temme K, Harrow AW, Kandala A, Chow JM, Gambetta JM. Supervised learning with quantum-enhanced feature spaces. Nature. 2019;567(7747):209–12. https://doi.org/10.1038/s41586-019-0980-2.
167. Hubregtsen T, Wierichs D, Gil-Fuster E, Derks P-JHS, Faehrmann PK, Meyer JJ. Training quantum embedding kernels on near-term quantum computers. Phys Rev A. 2022;106:042431. https://doi.org/10.1103/PhysRevA.106.042431.
168. Schuld M. Supervised quantum machine learning models are kernel methods. 2021. https://arxiv.org/abs/2101.11020.
169. Thanasilp S, Wang S, Cerezo M, et al. Exponential concentration in quantum kernel methods. Nat Commun. 2024;15:5200. https://doi.org/10.1038/s41467-024-49287-w.
170. Huang H-Y, Broughton M, Mohseni M, et al. Power of data in quantum machine learning. Nat Commun. 2021;12(1):2631. https://doi.org/10.1038/s41467-021-22539-9.
171. Miroszewski A, Nalepa J, Saux BL, Mielczarek J. Quantum machine learning for remote sensing: exploring potential and challenges. 2023. arXiv preprint arXiv:2311.07626. Submitted on 13 Nov 2023.
172. Miroszewski A, Mielczarek J, Czelusta G, Szczepanek F, Grabowski B, Le Saux B, Nalepa J. Detecting clouds in multispectral satellite images using quantum-kernel support vector machines. IEEE J Sel Top Appl Earth Obs Remote Sens. 2023;16:7601–13. https://doi.org/10.1109/JSTARS.2023.3304122.
173. Gupta MK, Beseda M, Gawron P. How quantum computing-friendly multispectral data can be? In: IGARSS 2022 - 2022 IEEE international geoscience and remote sensing symposium. 2022. p. 4153–6. https://doi.org/10.1109/IGARSS46834.2022.9883676.
174. Peters E, Caldeira J, Ho A, et al. Machine learning of high dimensional data on a noisy quantum processor. npj Quantum Inf. 2021;7:161. https://doi.org/10.1038/s41534-021-00498-9.
175. Lu C-Y, Cao Y, Peng C-Z, Pan J-W. Micius quantum experiments in space. Rev Mod Phys. 2022;94:035001. https://doi.org/10.1103/RevModPhys.94.035001.
176. Wei S, Jing B, Zhang X, Liao J, Yuan C, Fan B, Lyu C, Zhou D, Wang Y, Deng G, Song H, Oblak D, Guo G, Zhou Q. Towards real-world quantum networks: a review. Laser Photonics Rev. 2022;**16**(3). https://doi.org/10.1002/lpor.202100219.
177. Parny L, Alibart O, Debaud J, Gressani S, Lagarrigue A, Martin A, Metrat A, Schiavon M, Troisi T, Diamanti E, Gélard P, Kerstel E, Tanzilli S, Van Den Bossche M. Satellite-based quantum information networks: use cases, architecture, and roadmap. Commun Phys. 2023;**6**(1). https://doi.org/10.1038/s42005-022-01123-7.
178. Sidhu JS, Joshi SK, Gündoğan M, Brougham T, Lowndes D, Mazzarella L, Krutzik M, Mohapatra S, Dequal D, Vallone G, Villoresi P, Ling A, Jennewein T, Mohageg M, Rarity JG, Fuentes I, Pirandola S, Oi DKL. Advances in space quantum communications. IET Quantum Commun. 2021;2(4):182–217. https://doi.org/10.1049/qtc2.12015.
179. Bluvstein D, Evered SJ, Geim AA, Li SH, Zhou H, Manovitz T, Ebadi S, Cain M, Kalinowski M, Hangleiter D, et al. Logical quantum processor based on reconfigurable atom arrays. Nature. 2024;626(7997):58–65.
180. Cavaliere F, Prati E, Poti L, Muhammad I, Catuogno T. Secure quantum communication technologies and systems: from labs to markets. Quantum Rep. 2020;2(1):80–106.
181. Vasylyev D, Semenov AA, Vogel W, Günthner K, Thurn A, Bayraktar O, Marquardt C. Free-space quantum links under diverse weather conditions. Phys Rev A. 2017;96:043856. https://doi.org/10.1103/PhysRevA.96.043856.
182. Liorni C, Kampermann H, Bruß D. Satellite-based links for quantum key distribution: beam effects and weather dependence. New J Phys. 2019;21(9):093055. https://doi.org/10.1088/1367-2630/ab41a2.
183. Khatri S, Brady AJ, Desporte RA, Bart MP, Dowling JP. Spooky action at a global distance: analysis of space-based entanglement distribution for the quantum Internet. npj Quantum Inf. 2021;**7**(1). https://doi.org/10.1038/s41534-020-00327-5.
184. Chiti F, Picchi R, Pierucci L. A survey on non-terrestrial quantum networking: challenges and trends. Comput Netw. 2024;**110668**.
185. Wehner S, Elkouss D, Hanson R. Quantum Internet: a vision for the road ahead. Science. 2018;362(6412):9288. https://doi.org/10.1126/science.aam9288.
186. Caleffi M, Amoretti M, Ferrari D, Illiano J, Manzalini A, Cacciapuoti AS. Distributed quantum computing: a survey. Comput Netw. 2024;254:110672. https://doi.org/10.1016/j.comnet.2024.110672.
187. Cuomo D, Caleffi M, Cacciapuoti AS. Towards a distributed quantum computing ecosystem. IET Quantum Commun. 2020;1(1):3–8. https://doi.org/10.1049/iet-qtc.2020.0002.
188. Gomez AM, Patti TL, Anandkumar A, Yelin SF. Near-term distributed quantum computation using mean-field corrections and auxiliary qubits. Quantum Sci Technol. 2024;9(3):035022.
189. Kshemkalyani AD, Singhal M. Distributed computing: principles, algorithms, and systems. Cambridge: Cambridge University Press; 2008.
190. Awschalom DD, Hanson R, Wrachtrup J, Zhou BB. Quantum technologies with optically interfaced solid-state spins. Nat Photonics. 2018;12(9):516–27.

191. Couteau C, Barz S, Durt T, Gerrits T, Huwer J, Prevedel R, Rarity J, Shields A, Weihs G. Applications of single photons to quantum communication and computing. Nat Rev Phys. 2023;5(6):326–38.
192. Wolfowicz G, Heremans FJ, Anderson CP, Kanai S, Seo H, Gali A, Galli G, Awschalom DD. Quantum guidelines for solid-state spin defects. Nat Rev Mater. 2021;6(10):906–25.
193. Dréau A, Tchebotareva A, Mahdaoui AE, Bonato C, Hanson R. Quantum frequency conversion of single photons from a nitrogen-vacancy center in diamond to telecommunication wavelengths. Phys Rev Appl. 2018;9(6):064031.
194. Stolk A, Enden KL, Roehsner M-C, Teepe A, Faes SO, Bradley C, Cadot S, Rantwijk J, Te Raa I, Hagen R, et al. Telecom-band quantum interference of frequency-converted photons from remote detuned nv centers. PRX Quantum. 2022;3(2):020359.
195. Cacciapuoti AS, Illiano J, Viscardi M, Caleffi M. Multipartite entanglement distribution in the quantum Internet: knowing when to stop! IEEE Trans Netw Serv Manag. 2024;21(6):6041–58. https://doi.org/10.1109/tnsm.2024.3452326.
196. Ecker S, Bouchard F, Bulla L, Brandt F, Kohout O, Steinlechner F, Fickler R, Malik M, Guryanova Y, Ursin R, Huber M. Overcoming noise in entanglement distribution. Phys Rev X. 2019;**9**(4). https://doi.org/10.1103/physrevx.9.041042.
197. Daiss S, Langenfeld S, Welte S, Distante E, Thomas P, Hartung L, Morin O, Rempe G. A quantum-logic gate between distant quantum-network modules. Science. 2021;371(6529):614–7.
198. Noiri A, Takeda K, Nakajima T, Kobayashi T, Sammak A, Scappucci G, Tarucha S. A shuttling-based two-qubit logic gate for linking distant silicon quantum processors. Nat Commun. 2022;13(1):5740.
199. Cuomo D, Caleffi M, Krsulich K, Tramonto F, Agliardi G, Prati E, Cacciapuoti AS. Optimized compiler for distributed quantum computing. ACM Trans Quantum Comput. 2023;4(2):1–29.
200. Song J, Yang S, Liu P, Zhang H-L, Xue G-M, Mi Z-Y, Zhang W-G, Yan F, Jin Y-R, Yu H-F. Realization of high-fidelity perfect entangler between remote superconducting quantum processors. 2024. https://arxiv.org/abs/2407.20338.
201. Main D, Drmota P, Nadlinger DP, Ainley EM, Agrawal A, Nichol BC, Srinivas R, Araneda G, Lucas DM. Distributed quantum computing across an optical network link. Nature. 2025. https://doi.org/10.1038/s41586-024-08404-x.
202. Rubenok A, Pfister O, Ursin R. Quantum repeaters: from quantum networks to the quantum Internet. Quantum Netw, Phys Rep. 2023;775:14–30.
203. Fitzsimons JF. Private quantum computation: an introduction to blind quantum computing and related protocols. npj Quantum Inf. 2017;3(1):23.
204. Kimble HJ. The quantum Internet. Nature. 2008;453:1023–30. https://doi.org/10.1038/nature07127.
205. Monroe C, Kim J. Scaling the ion trap quantum processor. Science. 2013;339(6124):1164–9. https://doi.org/10.1126/science.1231298.
206. Devoret MH, Schoelkopf RJ. Superconducting circuits for quantum information: an outlook. Science. 2013;339(6124):1169–74. https://doi.org/10.1126/science.1231930.
207. Vepsäläinen AP, Karamlou AH, Orrell JL, Dogra AS, Loer B, Vasconcelos F, Kim DK, Melville AJ, Niedzielski BM, Yoder JL, et al. Impact of ionizing radiation on superconducting qubit coherence. Nature. 2020;584(7822):551–6.
208. Villoresi P, Jennewein T, Tamburini F, Aspelmeyer M, Bonato C, Ursin R, Pernechele C, Luceri V, Bianco G, Zeilinger A, et al. Experimental verification of the feasibility of a quantum channel between space and Earth. New J Phys. 2008;10(3):033038.
209. Vallone G, Bacco D, Dequal D, Gaiarin S, Luceri V, Bianco G, Villoresi P. Experimental satellite quantum communications. Phys Rev Lett. 2015;115(4):040502.
210. Yin J, et al. Satellite-based entanglement distribution over 1200 kilometers. Science. 2017;356(6343):1140–4. https://doi.org/10.1126/science.aan3211.
211. Kim J, Park J, Kim H-S, Kim G, Kim JT, Park J, Moon K, Kwak S-C, Kim M-S, Ju JJ. Fully controllable time-bin entangled states distributed over 100-km single-mode fibers. EPJ Quantum Technol. 2024;11(1):53.
212. Bulla L, et al. Nonlocal temporal interferometry for highly resilient free-space quantum communication. Phys Rev X. 2023;13(2):021001. https://doi.org/10.1103/PhysRevX.13.021001.
213. Erhard M, Krenn M, Zeilinger A. Advances in high-dimensional quantum entanglement. Nat Rev Phys. 2020;2(7):365–81. https://doi.org/10.1038/s42254-020-0193-5.
214. Zhou L, et al. Experimental quantum communication overcomes the rate-loss limit without global phase tracking. Phys Rev Lett. 2023;130(25):250801. https://doi.org/10.1103/PhysRevLett.130.250801.
215. Boaron A, et al. Secure quantum key distribution over 421 km of optical fiber. Phys Rev Lett. 2018;121(19):190502. https://doi.org/10.1103/PhysRevLett.121.190502.
216. Briegel H-J, Dür W, Cirac JI, Zoller P. Quantum repeaters: the role of imperfect local operations in quantum communication. Phys Rev Lett. 1998;81:5932–5. https://doi.org/10.1103/PhysRevLett.81.5932.
217. Żukowski M, Zeilinger A, Horne MA, Ekert AK. "Event-ready-detectors" Bell experiment via entanglement swapping. Phys Rev Lett. 1993;71:4287–90. https://doi.org/10.1103/PhysRevLett.71.4287.
218. Azuma K, Economou SE, Elkouss D, Hilaire P, Jiang L, Lo H-K, Tzitrin I. Quantum repeaters: from quantum networks to the quantum Internet. Rev Mod Phys. 2023;95:045006. https://doi.org/10.1103/RevModPhys.95.045006.
219. Heshami K, et al. Quantum memories: emerging applications and recent advances. J Mod Opt. 2016;63(20):2005–28. https://doi.org/10.1080/09500340.2016.1148212.
220. Mol J-M, Esguerra L, Meister M, Bruschi DE, Schell AW, Wolters J, Wörner L. Quantum memories for fundamental science in space. Quantum Sci Technol. 2023;8(2):024006.
221. Thomas SE, Wagner L, Joos R, Sittig R, Nawrath C, Burdekin P, Buy Wenniger IM, Rasiah MJ, Huber-Loyola T, Sagona-Stophel S, et al. Deterministic storage and retrieval of telecom light from a quantum dot single-photon source interfaced with an atomic quantum memory. Sci Adv. 2024;10(15):7346.
222. Chiti F, Fantacci R, Picchi R, Pierucci L. Mobile control plane design for quantum satellite backbones. IEEE Netw. 2022;36(1):91–7.
223. Ferrari D, Cacciapuoti AS, Amoretti M, Caleffi M. Compiler design for distributed quantum computing. IEEE Trans Quantum Eng. 2021;2:1–20.
224. Bennett CH, Brassard G. Quantum cryptography: public key distribution and coin tossing. In: Proceedings of IEEE international conference on computers, systems and signal processing. Bangalore, India. 1984. p. 175–9.
225. Ekert AK. Quantum cryptography based on Bell's theorem. Phys Rev Lett. 1991;67(6):661–3. https://doi.org/10.1103/PhysRevLett.67.661.

226. Pirandola S, et al. Advances in quantum cryptography. Adv Opt Photonics. 2020;12(4):1012–236. https://doi.org/10.1364/AOP.361502.
227. Neumann SP, Joshi SK, Fink M, Scheidl T, Blach R, Scharlemann C, Abouagaga S, Bambery D, Kerstel E, Barthelemy M, Ursin R. Q3sat: quantum communications uplink to a 3u cubesat—feasibility & design. EPJ Quantum Technol. 2018;**5**(1). https://doi.org/10.1140/epjqt/s40507-018-0068-1.
228. Oi DK, Ling A, Vallone G, Villoresi P, Greenland S, Kerr E, Macdonald M, Weinfurter H, Kuiper H, Charbon E, et al. Cubesat quantum communications mission. EPJ Quantum Technol. 2017;4:6.
229. Thompson R, Mahageg M, Kellogg J, Mosher T, Ramadorai R, Ruttley T, Wagner E, Woerner L. A quantum science space station. Quantum Sci Technol. 2022;8(1):014002.
230. Degen CL, Reinhard F, Cappellaro P. Quantum sensing. Rev Mod Phys. 2017;89:035002. https://doi.org/10.1103/RevModPhys.89.035002.
231. Marciniak CD, Feldker T, Pogorelov I, Kaubruegger R, Vasilyev DV, Bijnen R, Schindler P, Zoller P, Blatt R, Monz T. Optimal metrology with programmable quantum sensors. Nature. 2022;603(7902):604–9.
232. Cimini V, Valeri M, Piacentini S, Ceccarelli F, Corrielli G, Osellame R, Spagnolo N, Sciarrino F. Variational quantum algorithm for experimental photonic multiparameter estimation. npj Quantum Inf. 2024;10(1):26. https://doi.org/10.1038/s41534-024-00821-0.
233. Beckey JL, Cerezo M, Sone A, Coles PJ. Variational quantum algorithm for estimating the quantum Fisher information. Phys Rev Res. 2022;4(1):013083.
234. MacLellan B, Roztocki P, Czischek S, Melko RG. End-to-end variational quantum sensing. npj Quantum Inf. 2024;10(1):118. https://doi.org/10.1038/s41534-024-00914-w.
235. Strangfeld A, Carraz O, Eiden A, Heliere A, Silvestrin P. Quantum sensing for Earth observation at the European Space Agency: latest developments, challenges, and future prospects. Sen Syst Next-Gener Satell XXVII. 2023;12729:56–62.
236. Kantsepolsky B, Aviv I, Weitzfeld R, Bordo E. Exploring quantum sensing potential for systems applications. IEEE Access. 2023;11:31569–82. https://doi.org/10.1109/ACCESS.2023.3262506.
237. Alonso I, Alpigiani C, Altschul B, Araújo H, Arduini G, Arlt J, Badurina L, Balaž A, Bandarupally S, Barish BC, et al. Cold atoms in space: community workshop summary and proposed road-map. EPJ Quantum Technol. 2022;9(1):30.
238. Frye K, Abend S, Bartosch W, Bawamia A, Becker D, Blume H, Braxmaier C, Chiow S-W, Efremov MA, Ertmer W, et al. The Bose-Einstein condensate and cold atom laboratory. EPJ Quantum Technol. 2021;8(1):1.
239. Adams CS, Pritchard JD, Shaffer JP. Rydberg atom quantum technologies. J Phys, B At Mol Opt Phys. 2019;53(1):012002. https://doi.org/10.1088/1361-6455/ab52ef.
240. Schlossberger N, Prajapati N, Berweger S, Rotunno AP, Artusio-Glimpse AB, Simons MT, Sheikh AA, Norrgard EB, Eckel SP, Holloway CL. Rydberg states of alkali atoms in atomic vapour as si-traceable field probes and communications receivers. Nat Rev Phys. 2024;6(10):606–20. https://doi.org/10.1038/s42254-024-00756-7.
241. Barry JF, Schloss JM, Bauch E, Turner MJ, Hart CA, Pham LM, Walsworth RL. Sensitivity optimization for nv-diamond magnetometry. Rev Mod Phys. 2020;92:015004. https://doi.org/10.1103/RevModPhys.92.015004.
242. Gallego Torromé R, Barzanjeh S. Advances in quantum radar and quantum lidar. Prog Quantum Electron. 2024;93:100497. https://doi.org/10.1016/j.pquantelec.2023.100497.
243. Doherty MW, Manson NB, Delaney P, Jelezko F, Wrachtrup J, Hollenberg LCL. The nitrogen-vacancy colour centre in diamond. Phys Rep. 2013;528(1):1–45. https://doi.org/10.1016/j.physrep.2013.02.001.
244. Ruf M, Wan NH, Choi H, Englund D, Hanson R. Quantum networks based on color centers in diamond. J Appl Phys. 2021;130(7):070901. https://doi.org/10.1063/5.0056534 https://pubs.aip.org/aip/jap/article-pdf/doi/10.1063/5.0056534/20023329/070901_1_5.0056534.pdf.
245. Graham SM, Newman AJ, Stephen CJ, Edmonds AM, Twitchen DJ, Markham ML, Morley GW. On the road with a diamond magnetometer. Diam Relat Mater. 2025;**111945**. https://doi.org/10.1016/j.diamond.2025.111945.
246. Wang X, Li W, Moran B, Gibson BC, Hall LT, Simpson DA, Kealy AN, Greentree AD. Quantum diamond magnetometry for navigation in GNSS denied environments. In: Freymueller JT, Sánchez L, editors. Gravity, positioning and reference frames. Cham: Springer; 2024. p. 87–92.
247. Soshenko VV, Bolshedvorskii SV, Rubinas O, Sorokin VN, Smolyaninov AN, Vorobyov VV, Akimov AV. Nuclear spin gyroscope based on the nitrogen vacancy center in diamond. Phys Rev Lett. 2021;**126**(19). https://doi.org/10.1103/physrevlett.126.197702.
248. Beerden Y, Vandebosch R, Ermakova A, Nesladek M, Hruby J. Long-term stability assessment of quantum diamond magnetometers in low Earth orbit. In: IGARSS 2024 - 2024 IEEE international geoscience and remote sensing symposium. 2024. p. 465–8. https://doi.org/10.1109/IGARSS53475.2024.10641354.
249. Mandarino A, Torta P, Casati R, Prati E. Quantum computing for deep space physics applications. In: International astronautical congress. 75th international astronautical federation. Milan, Italy. 2024. p. 273–7. https://doi.org/10.52202/078361-0033.
250. Dailey C, Bradley C, Jackson Kimball DF, Sulai IA, Pustelny S, Wickenbrock A, Derevianko A. Quantum sensor networks as exotic field telescopes for multi-messenger astronomy. Nat Astron. 2021;5(2):150–8.
251. Panelli G, Roberts BM, Derevianko A. Applying the matched-filter technique to the search for dark matter transients with networks of quantum sensors. EPJ Quantum Technol. 2020;7(1):5.
252. Khabiboulline ET, Borregaard J, De Greve K, Lukin MD. Optical interferometry with quantum networks. Phys Rev Lett. 2019;123:070504. https://doi.org/10.1103/PhysRevLett.123.070504.
253. Tsang M, Nair R, Lu X-M. Quantum theory of superresolution for two incoherent optical point sources. Phys Rev X. 2016;6:031033. https://doi.org/10.1103/PhysRevX.6.031033.
254. Santamaria L, Sgobba F, Lupo C. Single-photon sub-Rayleigh precision measurements of a pair of incoherent sources of unequal intensity. Opt Quantum. 2024;2(1):46–56.
255. Huang Z, Lupo C. Quantum hypothesis testing for exoplanet detection. Phys Rev Lett. 2021;127:130502. https://doi.org/10.1103/PhysRevLett.127.130502.
256. Moll N, Barkoutsos P, Bishop LS, Chow JM, Cross A, Egger DJ, Filipp S, Fuhrer A, Gambetta JM, Ganzhorn M, et al. Quantum optimization using variational algorithms on near-term quantum devices. Quantum Sci Technol. 2018;3(3):030503.

257. Da Silva M, Ryan-Anderson C, Bello-Rivas J, Chernoguzov A, Dreiling J, Foltz C, Gaebler J, Gatterman T, Hayes D, Hewitt N, et al. Demonstration of logical qubits and repeated error correction with better-than-physical error rates. 2024. arXiv preprint arXiv:2404.02280.
258. Martinis JM. Saving superconducting quantum processors from decay and correlated errors generated by gamma and cosmic rays. npj Quantum Inf. 2021;7(1):90.
259. Wintersperger K, Dommert F, Ehmer T, Hoursanov A, Klepsch J, Mauerer W, Reuber G, Strohm T, Yin M, Luber S. Neutral atom quantum computing hardware: performance and end-user perspective. EPJ Quantum Technol. 2023;10(1):32.
260. Shi X-F. Single-site Rydberg addressing in 3d atomic arrays for quantum computing with neutral atoms. J Phys, B At Mol Opt Phys. 2020;53(5):054002.
261. LeCun Y, Bottou L, Bengio Y, Haffner P. Gradient-based learning applied to document recognition. Proc IEEE. 1998;86(11):2278–324.
262. Russakovsky O, Deng J, Su H, Krause J, Satheesh S, Ma S, Huang Z, Karpathy A, Khosla A, Bernstein M, et al. Imagenet large scale visual recognition challenge. Int J Comput Vis. 2015;115:211–52.
263. He K, Zhang X, Ren S, Sun J. Deep residual learning for image recognition. In: Proceedings of the IEEE conference on computer vision and pattern recognition. 2016. p. 770–8.
264. Gu J, Wang Z, Kuen J, Ma L, Shahroudy A, Shuai B, Liu T, Wang X, Wang G, Cai J, et al. Recent advances in convolutional neural networks. Pattern Recognit. 2018;77:354–77.

Publisher's note

Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.