


Quantum Algorithms Applied to Satellite Mission Planning for Earth Observation

Serge Rainjonneau, Igor Tokarev, Sergei Iudin, Saaketh Rayaprolu, Karan Pinto, Daria Lemtiuzhnikova, Miras Koblan, Egor Barashov, Mo Kordzanganeh , Markus Pflitsch, and Alexey Melnikov 

Abstract—Earth imaging satellites are a crucial part of our everyday lives that enable global tracking of industrial activities. Use cases span many applications, from weather forecasting to digital maps, carbon footprint tracking, and vegetation monitoring. However, there are limitations; satellites are difficult to manufacture, expensive to maintain, and tricky to launch into orbit. Therefore, satellites must be employed efficiently. This poses a challenge known as the satellite mission planning problem, which could be computationally prohibitive to solve on large scales. However, close-to-optimal algorithms, such as greedy reinforcement learning and optimization algorithms can often provide satisfactory resolutions. This article introduces a set of quantum algorithms to solve the mission planning problem and demonstrate an advantage over the classical algorithms implemented thus far. The problem is formulated as maximizing the number of high-priority tasks completed on real datasets containing thousands of tasks and multiple satellites. This work demonstrates that through solution-chaining and clustering, optimization and machine learning algorithms offer the greatest potential for optimal solutions. This article notably illustrates that a hybridized quantum-enhanced reinforcement learning agent can achieve a completion percentage of 98.5% over high-priority tasks, significantly improving over the baseline greedy methods with a completion rate of 75.8%. The results presented in this work pave the way to quantum-enabled solutions in the space industry and, more generally, future mission planning problems across industries.

Index Terms—Earth observation, quantum algorithms, quantum optimization, quantum reinforcement learning, satellite mission planning.

I. INTRODUCTION

THE reliable functioning of Earth-orbiting satellites crucially affects our everyday services such as connectivity [1], navigation [2], and media [3]. Most satellites receive dynamic instructions on executing their mission in orbit, and planning the exact sequence of tasks is critical to the efficiency and sustainability of the project [4]. Algorithmic optimization

solutions have been suggested [5], [6], [7] as a remedy to the planning problem. With the rise of quantum technologies, there is a need to explore how quantum computing can improve the time complexity or the quality of these solutions. This work focuses on planning the mission of Earth-orbiting imaging satellites using quantum machine learning [8], [9], [10] and optimization. Specifically, it explores using near-term quantum technologies to improve the solution's effectiveness today. Similar approaches, such as the contribution of Wu et al. [11], suggested an algorithm to overcome the scheduling task using quantum annealers. Stollenwerk et al. [12] reviewed the literature and found that although innovative developments existed, none revealed any practical advantage over classical approaches. The current work explores the interplay between classical and gate-based quantum computing algorithms. The practical advantage of employing this hybrid approach was shown in our earlier contributions [13], [14], [15].

In general, space mission planning can be a computationally hard problem [4] to solve; in large-scale missions, the size of the problem requires a prohibitive amount of computational resources. Finding an efficient plan requires the optimization of movement and the reordering of tasks to maximize the number of total completed tasks. In this work, each task is a request made to the satellite to capture an image of the surface of the Earth, and the aim is to maximize the number of images taken, given a list of all requests and a total available time. In this work, the imaging satellites orbit the Earth exactly on the Earth's terminator, which is the line separating Earth's sun-lit areas from the dark ones. In 24 hours, the satellite orbits approximately 15 times around the Earth, which leads to 15 orbits shown schematically in Fig. 3(a). To capture the image, the satellite must continuously aim the camera at the target area for a time known as the acquisition slot. Each requested image has an allocated data-take opportunity (DTO) window. The satellite must point in the appropriate direction within this window for the entire acquisition period to accomplish a request. To aim at an area on Earth, the satellite needs to rotate its camera to point in the desired direction. The latter movement introduces a time delay that, when added to the acquisition time, could limit the overall agility of the satellite in covering multiple areas. The efficient satellite mission plan will be able to choose the order of the acquisition requests to maximize the number of completed requests.

This article discusses the benefits of the optimization and reinforcement learning (RL) algorithms over a greedy baseline

Manuscript received 30 December 2022; revised 27 April 2023; accepted 30 May 2023. Date of publication 19 June 2023; date of current version 4 August 2023. (Corresponding author: Alexey Melnikov.)

Serge Rainjonneau is with the Thales Alenia Space, 31100 Toulouse, France (e-mail: serge.rainjonneau@thalesaleniaspace.com).

Igor Tokarev, Sergei Iudin, Saaketh Rayaprolu, Karan Pinto, Daria Lemtiuzhnikova, Miras Koblan, Egor Barashov, Mo Kordzanganeh, Markus Pflitsch, and Alexey Melnikov are with the Terra Quantum AG, 9000 St. Gallen, Switzerland (e-mail: igto@terraquantum.swiss; sj@terraquantum.swiss; sr@terraquantum.swiss; karan@terraquantum.swiss; dl@terraquantum.swiss; mik@terraquantum.swiss; eb@terraquantum.swiss; mko@terraquantum.swiss; pflitsch@terraquantum.swiss; ame@terraquantum.swiss).

Digital Object Identifier 10.1109/JSTARS.2023.3287154

algorithm, the unique advantage offered by the potential of quantum computing, and a demonstration of how quantum methods can be applied to improve machine learning and optimization models. Specifically, the novel quantum RL approach shown in Section III-D4 offers a completion rate of 98.5% on highest priority requests in a multisatellite system. The model hybridizes the AlphaZero approach in [16] and is trained on the QMware quantum cloud [17]. Section II introduces the practical problem setup in more detail, including the details on the data formatting in Section II-A, reordering the requests in Section II-B, relaying algorithm in Section II-C, and on solution chaining in Section II-D. Section III offers three algorithms and their respective results: Section III-B establishes a baseline greedy algorithm, and Sections III-C and III-D discuss the optimization and RL algorithms. Finally, Section IV summarizes the results.

II. SATELLITE MISSION PLANNING

Each satellite's orbit is constrained to the Earth's terminator. Each satellite can only rotate at a maximum of one degree per second to orient itself appropriately to capture images within the acquisition window. Furthermore, the DTO duration for each area to be captured is defined by the arc ranging within 45° (referred to as the depointing angle) from the apex point, which is the point directly above the center of the request on Earth. As a result of the width of these arcs, multiple requests can have substantial overlap in their acquisition windows. Furthermore, while coordinates for image requests are provided in latitude and longitude, the satellite coordinates support an Earth-centric inertial (ECI) format, providing more spatial location support. Finally, two datasets are tested in this work, including a single-satellite set containing 462 requests and a two-satellite system with 2000 requests. The main results of this work focus on the performance of various algorithms on the latter dataset and only for the requests with the highest priority.

A. Data Formatting

Two information sets were used in the preparation of this article:

- 1) information about the satellite motion, including
 - a) the orbit number;
 - b) time stamp; and
 - c) satellite position and velocity in Cartesian ECI coordinates; and
- 2) information about acquisition request, including
 - a) request ID;
 - b) request priority ranging from 1 to 4, where 4 denotes the lowest priority;
 - c) start and end times of the DTO window of the request;
 - d) the coordinates of the start and end of the median line;
 - e) satellite ID; and
 - f) Boolean values indicating the progress of the acquisition.

Fig. 1 demonstrates the satellite movements during acquisition and the data that must be tracked during the capture, such as the acquisition angles, the points at which the DTO begins and ends, and the median coordinates.

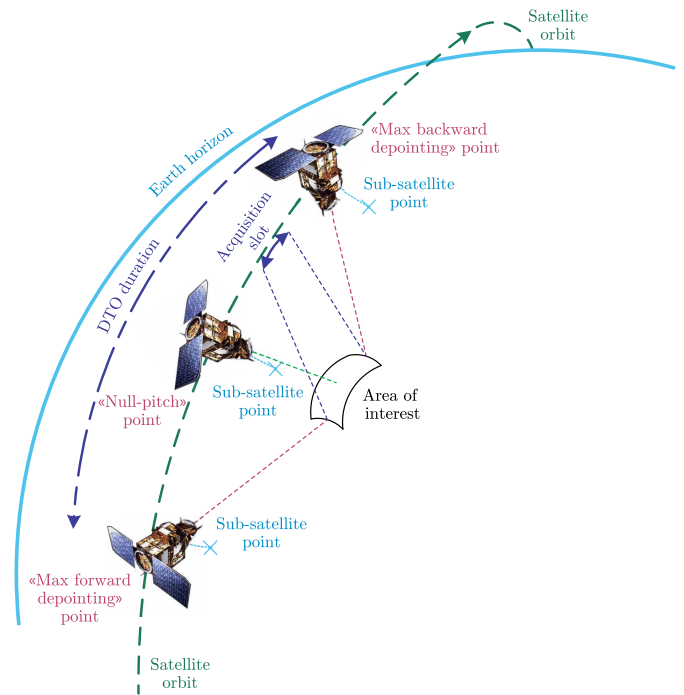


Fig. 1. Trajectory of a satellite orbiting on the Earth's terminator and across a DTO window. At the first and third positions on the orbital line, the satellite is at the ends of the request's DTO window, as its depointing angle is maximized at 45° . At its second position, it is at the apex point, directly above the request location.

B. Request Priority Ordering

Considering the priority associated with each request, it is possible to express this objective as maximizing the completed requests in the order of priority; the highest priority requests are considered first, and only the lower priority requests are considered upon completion. The exact quantification compares the number of higher priority requests accomplished. It moves on to the next priority in case of equality, but this work judges the algorithms by their completion rate performance purely on the high-priority π_1 request.

C. Relaying Algorithm

The satellite must rotate according to the relaying algorithm to move from one request to another. To compute this rotation, two points of interest are the final median points of the first request and the initial median point of the second request. These points, respectively, signify the time at which the first request was completed and the time at which the next request will begin. For a given transition, the relaying algorithm operates in two steps: 1) it computes the relative positions of each median point for the satellite, and 2) calculates in degrees the angle between these vectors. Assuming that the satellite rotates at a constant speed of one degree every second and that the Earth is a perfect sphere, the resultant angle can be approximated as the total relaying time in seconds.

D. Solution Chaining

Solution chaining considers the DTO windows, overlaps, and the windows' length compared to an acquisition time estimate.

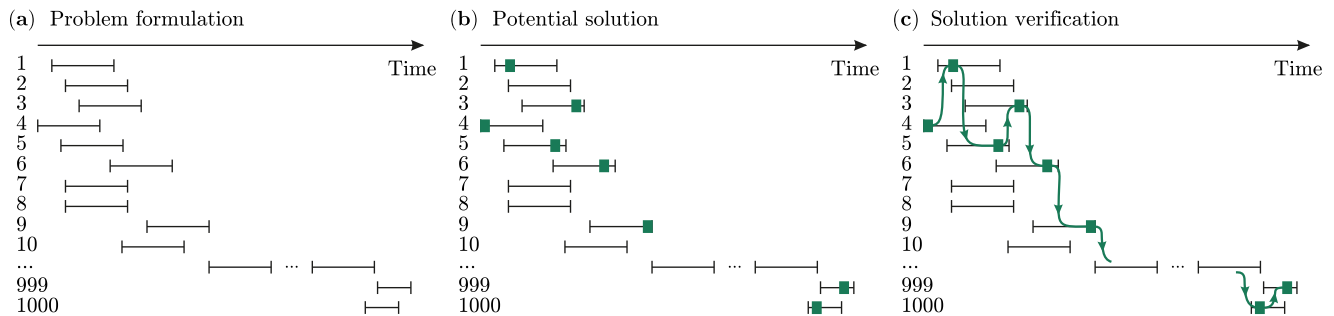


Fig. 2. (a) Visualization of DTO windows (horizontal bars) over a horizontal axis of time and a vertical axis indicating request index. (b) Visualization of solutions (green rectangles). (c) Chaining of requests based on their DTO windows and completion windows. As is evident, chaining provides a roadmap to connect all completed requests and map out the order of requests.

DTO windows and possible solutions within those windows are then visualized over an axis of time—see Fig. 2(b). The most efficient preliminary method for this setting is to sort the requests by the start of their DTO windows and map the acquisitions.

Next, the individual requests *chained*: the intermittent time between captures is calculated using the relaying algorithm, after which the algorithm ensures sufficient time for the relaying movement between the completion of each pair of consecutive requests, ensuring that a sequence of requests can be executed in the provided time. Chaining enables the suggested solutions within the DTO windows of each request, which are otherwise independently scattered, to be connected together into one comprehensive solution for satellite movement. Shown in Fig. 2(c) is a visualization of the chaining process between multiple requests over time.

III. RESULTS: QUANTUM ALGORITHMS FOR SATELLITE MISSION PLANNING

A. Clustering for Preprocessing

To deal with increasing complexity in problems is to reduce the size of the dataset and deal with fewer parameters or data points at once. Therefore, clustering is useful to stratify the data by similarities and shorten the calculations necessary for the overall program. The simplest method to cluster the dataset is sorting the samples based on a single feature of the data, a method known as bunching. The bunching algorithms are explained in Appendix A, but this section focuses on unsupervised-learning clustering methods such as the K-means algorithm [18], [19].

This algorithm is a well-known clustering algorithm, and its usage of physical distances in creating clusters makes it a natural fit for the Space Mission Planning problem. This algorithm is iterative, meaning it runs multiple iterations of the same steps and finally converges at a solution. The first step of this algorithm is the assignment stage; once k random points are initialized to represent the k cluster *centroids*, each point in the dataset is assigned to the closest centroid. Once this is complete, the values of all points assigned to each centroid k_n are averaged, and the value of k_n is updated to the newly obtained mean value. Once these two steps are completed, they are repeated until the centroids no longer shift between two iterations (generally to

TABLE I
RESULTS FOR THE GREEDY ALGORITHM BY PRIORITY, RUN OVER TWO SEPARATE DATASETS

Priority	1sat/462req	2sat/2000req
π_1	99.2%	75.8%
π_2	88.6%	36.1%
π_3	80.7%	19.2%
π_4	74.0%	18.2%

an error threshold), which is when the final clustering for the dataset is attained. The K-means algorithm clustered the dataset based on all the features of the requests: the DTO start and end times and the coordinates of the start and end of the median line. This ensured that clusters bunch the requests based on their DTO windows and geographical locations. The latter is crucial in efficiently using the relaying algorithm, as the geographically close requests are simpler to internavigate.

B. Greedy Algorithm

The classical greedy algorithm is used to provide a reference solution. First, the orbital and request data are separated by satellite, and then the requests are clustered. In each cluster, the algorithm begins at the satellite's given start time and increments by one second until it enters the DTO window of the first request. Once the time stamp enters a DTO window, the algorithm checks for enough time to complete a request. If so, it completes the request, increments the time by the sum of the relaying time and the acquisition time, adds the request ID to a list of completed IDs and moves on to the next request. The algorithm chooses the highest priority request if multiple requests are available at a timestamp. If insufficient time to complete a request, it simply moves on to the next request and repeats the process. Once the final request is completed, or the timestamp exceeds the final DTO window, the algorithm outputs the percentage of completed requests of each priority.

An important limitation of the greedy algorithm is its bandwidth for anticipation, but it provides a baseline for comparison. Its simplicity allows it to run faster than other models, making it a good benchmark for data preprocessing. The results in Fig. 3(b) and Table I show the greedy algorithm struggles as the complexity of the input data increases. Therefore, to solve this

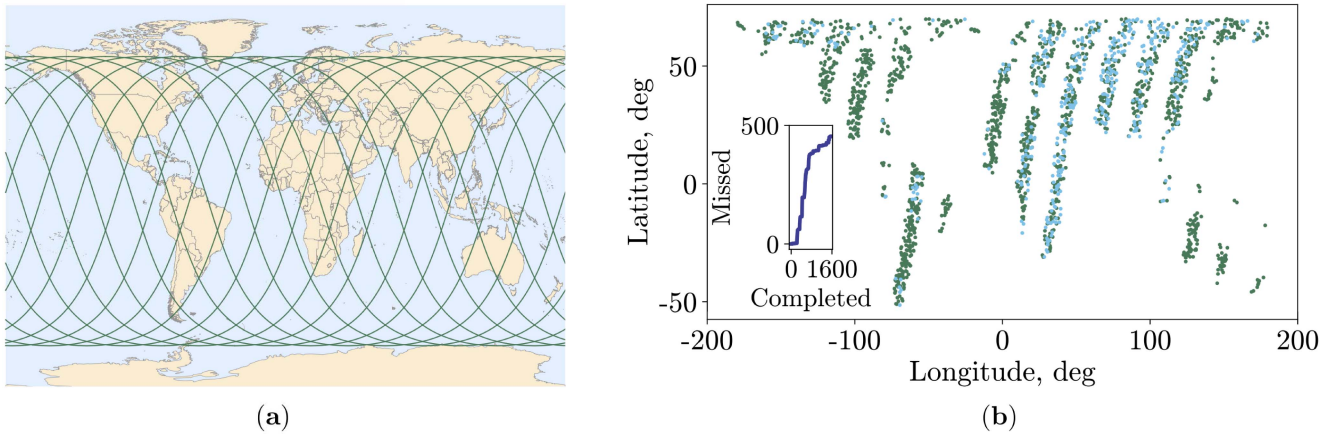


Fig. 3. (a) Schematic view of 15 satellite orbits performed within approximately 24 h. (b) Results of the greedy algorithm are shown on the map. The task contains 2000 requests that are planned for acquisition using two satellites. Green points correspond to completed requests, whereas blue points are missed requests. (Inset) Progression of completed compared to missed requests.

TABLE II
RESULTS OF THE GREEDY ALGORITHM USING CLASSICAL
K-MEANS CLUSTERING

Priority	KMeans + Greedy	
	1sat/462	2sat/2000req
π_1	99.1%	63.6%

problem, algorithms that can handle high levels of complexity without compromising accuracy are required. The results of the greedy algorithm serve as the baseline algorithm in this work. As evident in Tables I and II, the greedy algorithm solves the 1sat/462req dataset well in clustered and unclustered cases, but it struggles with the 2sat/2000req dataset. Thus, the latter was chosen as the comparison dataset as it allowed more possibilities for improvement. Furthermore, the unclustered greedy algorithm performs best on this dataset, so it was chosen as the baseline algorithm here, it should be noted that only the top priority π_1 requests were taken as the baseline, so the models' accuracy performance will be compared with the baseline π_1 2sat/2000req of 75.8%.

C. Optimization Methods

As discussed earlier, one potential way to improve the runtime of an algorithm was to break the data into smaller clusters to reduce the necessary processing power. However, attempting to build a fundamentally more powerful algorithm, such as an optimization model, could be more fruitful. Optimization problems involve information and formulations, including graphs, movement patterns and permutations, and multiple viable solutions, out of which one ideal solution must be determined by the metrics and constraints of the problem. As such, optimization methods are well suited for mission planning problems akin to the one at hand, which aims to map out the best possible course of action for a system of satellites seeking to maximize the number of completed requests. Optimization, however, is expensive in terms of time and computation; consequently, the potential of harnessing the quantum advantage for optimization problems

could be especially valuable in discovering and delivering a large speedup for mission planning problems.

Some research for optimization for space mission planning currently exists. For instance, in [20], the problem of Earth observation from a satellite (EOS) is investigated, focusing on obtaining images of certain areas of the Earth's surface related to customer requests. An optimization approach for EOS's daily photo selection (DPSP) is proposed. DPSP is related to operational management and planning processes, where each photo the client orders brings profit. Still, not all requests can be satisfied due to physical and technological limitations. The objective is, therefore, to select a subset of queries for which the profit is maximized, and the proposed algorithm is based on the metaheuristic ant colony optimization algorithm (ACO). Examples based on real data are used as reference problems. The calculations show that the proposed algorithm can generate competitive and promising solutions.

The most well-known quantum-friendly optimization method is the quadratic unconstrained binary optimization (QUBO) model, which encapsulates the set of all optimization problems with the following attributes:

- 1) all variables are treated as binary objects;
- 2) constraints, while not necessarily absolute, are enforced accordingly by use of the reward function.

This work explores the QUBO-adaptable formulation known as the *integer optimization model* for the space mission planning problem. This solution has benefits and drawbacks, but can serve as a viable solution to the satellite optimization problem.

The integer optimization model is conceptualized by encoding the desirable outcome of the solution in the cost function, which should be maximized or minimized on the lattice of integer points of the special feasible subset in the multidimensional space. This feasible subspace is defined by adding the constraints of the satellite mission planning problem. One of the most powerful approaches to that problem is the *branch and cut* algorithm [21], used in solvers. In addition, the following functions are incorporated to streamline the calculations for the algorithm to stay within the constraints. Incorporating these functions is

crucial to developing viable solutions as these functions, built over the Orekit space flight dynamics library [22], are the key to accurate space mission planning.

- 1) *Attitude Pointing*: given the satellite, timestamp, and a pair of longitude and latitude coordinates, this function returns the attitude (roll, pitch, and yaw) angles of the satellite when it is pointed toward the provided coordinates.
- 2) *Acquisition Duration*: given the starting and ending median coordinates of any request, this function returns the amount of time, in milliseconds, that the satellite would take to acquire the image.
- 3) *Maneuver Duration*: given the beginning and ending attitude angles (roll, pitch, and yaw), this function returns the amount of time necessary for the satellite to complete the maneuver from one attitude to the other.
- 4) *Read Ephemeris*: given a JSON file and a satellite ID, this function returns the satellite orbital information (directional speed, directional position, timestamps, etc.) in the form of an Orekit Ephemeris object, which is then used to perform calculations for other functions using the Orekit library.

Upon incorporating the above functions, revision of the constraints on this integer-based model, and modification of the clustering algorithm to generate a greater number of smaller clusters, the algorithm's runtime was maintained despite the significant increase in complexity. Given a set of requests, the variables central to this model include the priority of each request, the acquisition duration of each request, the start and the end of the DTO windows, and the location of the median points of the requests. From this, the algorithm calculates for each reasonable pair of the requests inside the small cluster all possible start times for the first request in a way that allows the relaying maneuver to the 2d request. This is achieved through simple iteration, using the Orekit functions described above: an initial value for the relaying maneuver duration is set to one second. Then, the maneuver start and end angles are calculated with the attitude pointing function. Afterward, suppose the rotation time between the obtained angles via the maneuver duration function is not equal to or greater than one second, and the DTO limitations are not violated. In that case, meaning the relaying time is equal to one second. Otherwise, the duration is increased by one second, and this procedure is repeated until either the duration is appropriate or the DTO is violated. The result of these computations t_{\min} is treated as the minimum relaying maneuver time.

Given a set of requests $f \in \mathbf{F}$, the priority of each request is denoted as π_f , and the acquisition duration as τ_f . The start of the DTO window (release time) for request f is represented as r_f and the deadline of the DTO window is represented as d_f . In addition, $\mathbf{Q} \subset \mathbb{N}$ indicates the order of accepted requests. As a result, index $q \in \mathbf{Q} = \{0, \dots, Q-1\}$ demonstrates the position in the queue of requests in which request f is accepted. The discretized rotation of the satellite at the inception of the acquisition of request f is indicated by $\alpha_f^{\text{start}} \in \mathfrak{A}_f^{\text{start}}$. Similarly, $\alpha_f^{\text{end}} \in \mathfrak{A}_f^{\text{end}}$ is the discretized rotation of the satellite after the acquisition of the request. However, all possible inceptions $b_{f\alpha}$ for requests can be streamlined in order of increasing time. Note

that $b_{f\alpha}$ is the set of all possible points in time at which it is possible to complete the f_1 request and then move to the f_2 request. That is, it is at a time greater than r_f . Thus, all the possible angles will correspond to these moments. The set of possible pairs is defined as \mathbf{L} , such that the maneuver $f_1 \rightarrow f_2$ is possible; i.e.,

$$\begin{aligned} \exists b_{f_1\alpha_{f_1}}, b_{f_2\alpha_{f_2}} : b_{f_1\alpha_{f_1}} + t_{\min}(b_{f_1\alpha_{f_1}} + \tau_{f_1}, f_1, f_2) \\ + \tau_{f_1} = b_{f_2\alpha_{f_2}} \end{aligned} \quad (1)$$

where $t_{\min}(t, f_1, f_2)$ is the minimum amount of time required for a satellite to rotate from its position at the end of f_1 at the moment t to the starting point of f_2 . For all $(f_1, f_2) \in \mathbf{L}$, the sets $\mathfrak{B}_{f_1f_2}^{\text{start}} \subseteq \mathfrak{A}_{f_1}^{\text{end}}$ and $\mathfrak{B}_{f_1f_2}^{\text{end}} \subseteq \mathfrak{A}_{f_2}^{\text{start}}$ and mapping $M_{f_1, f_2} : \mathfrak{B}_{f_1f_2}^{\text{start}} \rightarrow \mathfrak{B}_{f_1f_2}^{\text{end}}$ are defined such that $M_{f_1, f_2}(\alpha_1) = \alpha_2$, if $b_{f_1\alpha_1} + t_{\min}(b_{f_1\alpha_1} + \tau_{f_1}, f_1, f_2) + \tau_{f_1} = b_{f_2\alpha_2}$.

The binary variable x_{fq} is introduced such that

$$x_{fq} = \begin{cases} 1, & \text{if the } f\text{th request is started in the } q\text{th slot} \\ 0, & \text{otherwise.} \end{cases} \quad (2)$$

In addition, the variable $y_{f\alpha}$ is also introduced, where

$$y_{f\alpha} = \begin{cases} 1, & \text{if angle } \alpha \text{ is the start angle for request } f \\ 0, & \text{otherwise.} \end{cases} \quad (3)$$

Finally, the indicator variable $\kappa_{f_1f_2}$ shows if requests were completed successively one after another

$$\kappa_{f_1f_2} = \begin{cases} 1, & \text{if the } f_2\text{th follows right after the } f_1\text{th} \\ 0, & \text{else.} \end{cases} \quad (4)$$

The cost function is defined as

$$\sum_{f \in \mathbf{F}} \left(\sum_{q \in \mathbf{Q}} J_f x_{fq} - \sum_{\alpha \in \mathfrak{A}_f^{\text{start}}} \gamma_{f\alpha} y_{f\alpha} \right) \rightarrow \max \quad (5)$$

where J_f is the weight of request f , and $\gamma_{f\alpha}$ is a coefficient needed for penalizing time consuming solutions. $J_f = 1$ was selected for the lowest priority requests in each cluster. For the higher priority requests, the weight is greater than the sum of all lower priority requests' weights by one since each higher priority request is valued more than any amount of lower priority requests. The coefficients $\gamma_{f\alpha}$ range from 0 to $\frac{1}{Q}$ for the following definition:

$$\gamma_{f\alpha} = \frac{b_{f\alpha} - r_f}{Q(d_f - r_f - \tau_f + 1)}. \quad (6)$$

It guarantees that $\sum_f \sum_{\alpha \in \mathfrak{A}_f^{\text{start}}} \gamma_{f\alpha} y_{f\alpha} < 1$, and consequently, the completion of the lowest priority request will be more important than the particular order of requests, but the earliest possible completion of each request is preferable.

As was the case with the other optimization model, more than one request cannot be executed in the same order, and each request should be completed not more than once

$$\forall f \in \mathbf{F}, \sum_{q \in \mathbf{Q}} x_{fq} \leq 1 \quad (7)$$

$$\forall q \in \mathbf{Q}, \sum_{f \in \mathbf{F}} x_{fq} \leq 1. \quad (8)$$

Any completed request is started with one particular possible angle

$$\sum_{\alpha \in \mathfrak{A}_f^{\text{start}}} y_{f\alpha} = \sum_{q \in \mathbf{Q}} x_{fq} \quad \forall f \in \mathbf{F}. \quad (9)$$

All requests are completed one after another without empty slots in line until the satellite stops

$$\sum_{f_1 \in \mathbf{F}} x_{f_1(q-1)} \geq \sum_{f_2 \in \mathbf{F}} x_{f_2q} \quad \forall q > 0. \quad (10)$$

To evaluate the variable $\kappa_{f_1 f_2}$ appropriately, the following system of linear equations is introduced, and each excludes the impossible values of $\kappa_{f_1 f_2}$. First, if requests f_1 and f_2 were completed straight one after another, then $\kappa_{f_1, f_2} = 1$:

$$\kappa_{f_1 f_2} + 1 \geq x_{f_1, q-1} + x_{f_2, q} \quad \forall f_1, f_2 \in \mathbf{F}, q > 0. \quad (11)$$

Each request is followed by not more than one request, and each request follows after not more than one request

$$\sum_{f_1 \in \mathbf{F}} \kappa_{f_1, f_2} \leq 1 \quad \forall f_2 \in \mathbf{F} \quad (12)$$

$$\sum_{f_2 \in \mathbf{F}} \kappa_{f_1, f_2} \leq 1 \quad \forall f_1 \in \mathbf{F}. \quad (13)$$

If one request follows another one, then each of these requests was completed in some order

$$\kappa_{f_1 f_2} \leq \sum_{q > 0} x_{f_2, q} \quad \forall f_1, f_2 \in \mathbf{F} \quad (14)$$

$$\kappa_{f_1 f_2} \leq \sum_{q < Q-1} x_{f_1, q} \quad \forall f_1, f_2 \in \mathbf{F}. \quad (15)$$

With the use of predefined mapping M_{f_1, f_2} , if the request f_1 starts with angle α_1 , and the request f_2 is the next one, then the next acquisition start angle is fixed

$$y_{f_2 \alpha_2} + 1 \geq \kappa_{f_1, f_2} + y_{f_1 \alpha_1} \quad \forall \alpha_1 \in \mathfrak{B}_{f_1 f_2}^{\text{start}} \quad (16)$$

where $\alpha_2 = M_{f_1, f_2}(\alpha_1)$.

In addition, maneuver $f_1 \rightarrow f_2$ is possible only with particular initial angles for f_1 . Otherwise, the satellite will not have enough time to finish the acquisition of the request f_1 and move to f_2

$$\sum_{\alpha_1 \in \mathfrak{B}_{f_1 f_2}^{\text{start}}} y_{f_1 \alpha_1} \geq \kappa_{f_1 f_2} \quad \forall (f_1, f_2) \in \mathbf{L}. \quad (17)$$

On the other hand, if $f_1 \rightarrow f_2$ is an impossible transaction, then

$$\kappa_{f_1, f_2} = 0 \quad \forall (f_1, f_2) \notin \mathbf{L}. \quad (18)$$

The final constraint fixes the acquisition start angle for the first request in a queue as the earliest possible angle

$$y_{f \alpha_0} \geq x_{f0} \quad \forall f \in \mathbf{F} \quad (19)$$

where α_0 is an angle corresponding to the beginning of the DTO for request f .

In contrast to the greedy algorithm, which considered the requests in the order of open DTO windows, this model is more intelligent in finding the best path to fit the maximal number of requests. In the end, the π_1 requests reached 98.1% completion using the Gurobi solver [23]. The solution obtained through optimization is partly depicted in Fig. 4. Furthermore, the clusters are connected, calculating the minimum relaying maneuver time from the last request of the previous cluster to that of the next. This is illustrated in Fig. 4 as cutting the beginning of every DTO from the start if it proves impossible to rotate toward the request in this period. After this procedure, the next cluster can be treated as independent.

One of the greatest benefits of this model is its compatibility with both near-term and long-term quantum technology. As a linear optimization model that uses a grid of binary parameters, it can be transformed into QUBO, as it is shown in Appendix B, and fit rather quickly to the quantum Ising model, a model containing arrays of qubits in a grid, where their spin states depend on their neighbors. Furthermore, as this model functions akin to a minimization problem, it would be extremely efficient to run on a quantum annealing machine [24], which could solve optimization problems by slightly changing the Hamiltonian from a given initial state with a known minimum to a new state, representing the optimal solution. However, it is complicated to use the satellite mission planning problem in the QUBO form via currently available classical or quantum devices. For example, both D-Wave's Leap Hybrid solver [25] and Gurobi have difficulties solving even a small cluster with four requests over a time limit of 1 min. Still, linear programming can obtain the solution for the 2000 requests and two satellites dataset in less than 3 min. This runtime was achieved by decomposing the clusters so that their final size was small enough to achieve approximately the same time as the greedy algorithm. Note that if the same number of requests were simultaneously considered, the runtime would be noticeably longer.

D. Reinforcement Learning

RL is a machine learning paradigm in which an agent interacts with some environment and trains through informed trial and error. The RL training algorithm uses reward functions to assign value to the agent's actions in any state of the environment. Generally, a state can have constraints and features. The RL agent can use a policy model to decide on an action given a state, which subsequently affects the environment and transforms the state. Suppose the resultant state of the environment is engineered by the data scientist to contain a positive reward. In that case, the policy model is trained to take the appropriate action to maximize the probability of achieving that reward. The Environment is a function of a triplet of variables (S, A, P) , where S is a state space, A is an action space, and P is a transition function. When the reward function r is factored in, a Markov decision process (MDP) is generated with the property (S, A, P, r) , $r : S \times A \rightarrow \mathbb{R}$. The Agent starts from state s_0 and takes action a_0 , for which the reward r_0 is obtained in each step of training and subsequently trains by producing trajectories $T := (s_0, a_0, r_0, s_1, a_1, r_1, s_2, a_2, r_2, \dots)$.




Fig. 4. Chaining results of classical integer optimization algorithm within clusters of data points.

1) *RL Environments*: Two environments are developed for this AEOS task: satellite- and request-centred environments. For each environment, requests are sorted by their DTO open time.

The satellite-centred environment: The agent views the problem from the satellite's perspective and considers the 100 closest (by DTO open time) data points for each satellite, each utilizing ten data parameters. Moreover, three additional features were also made available to the agent: the current time for the satellite and the latitude and longitude of the starting point of the last completed request, creating a total observation space of size 1003. The number of nearest requests is a variable that depends on task size. A Boolean flag is one of the features kept for each request, marking each data point as either complete or incomplete to avoid redundancy, and the observation space is made to only consider requests that are marked as incomplete to expedite the computations required by the agent, which is built as a neural network.

Once the request is selected, the agent is rewarded with 1 if that request is completed and 0 otherwise. With this process recurring, the agent attempts to complete as many requests as possible until the time of the satellite is greater than the DTO windows of all remaining requests, meaning they can no longer be completed. The agent works with each satellite and predicts which request must be done. The data used in this environment were artificially generated. *The request-centred environment*: at each step, the agent views the problem from the perspective of a request, deciding which satellite is best suited to complete it. The 5 nearest request options are determined by DTO open time for each satellite, and the request to execute is chosen by the minimum request execution time. This minimum time is calculated as the sum of the satellite timestamp, maneuver duration, and the acquisition time due to solution chaining. It must be less than the request DTO end time for completion. This procedure is then iterated with the next batch of five nearest satellites.

2) *Proximal Policy Optimization*: The proximal policy optimization (PPO) algorithm, shown in Fig. 5 and first introduced in [26], was implemented to provide a mission planning policy. In RL, a policy is an operation that maps an action space to a state space. The agent learns the best action for each situation by calculating the policy gradients. In other words, the agent uses gradient descent to calculate the expected value of each action at a certain state space and determine which action has the likelihood of the highest reward. The equation for the PPO

algorithm is as follows:

$$L(\theta) = \bar{E}_t[\min(r_t(\theta)\bar{A}_t, \text{clip}(r_t(\theta), 1 - \epsilon, 1 + \epsilon)\bar{A}_t)] \quad (20)$$

where \bar{E}_t is the current expected value of the policy, θ is the policy parameter, ϵ is the hyperparameter, \bar{A}_t is the estimated advantage provided by the PPO at time t , and $r_t(\theta)$ is the importance sampling ratio. This ratio, derived from the Monte Carlo sampling methods [26], denotes the ratio of probabilities under both the old and the new policies. By keeping the value of ϵ small, the model ensures that the update on the policy at each increment is not too large; as a consequence, the learning done by the model stays relevant.

3) *Hybrid-Quantum PPO*: Inspired by the quantum advantage shown in [13], [14], [15], [27], [28], [29], [30], this section investigates the utility of a hybrid-quantum neural network as a policy model for RL. Fig. 5(b) illustrates the specific policy network within a hybrid MLP model. In the quantum-classical network, the output of the classical neural network is used as inputs to a parametrized quantum circuit (PQC). PQCs are quantum circuits that use parametrized quantum gates [31], [32] such as Pauli rotations to encode data x and trainable parameters θ . The four-qubit PQC used in this work consists of three significant components: variational, data encoding, and measurement. The variational layer comprises a layer of Pauli-X rotations encoding four trainable parameters and ring-shaped CNOTs for entanglement. The data encoding layer embedded four features in parallel using Pauli-Z rotations, and the measurement layer comprised two Z-basis measurements on the first two qubits. In sequence, four qubits were initialized in the ground state, and then a variational layer with randomly initialized parameters was appended, followed by eight repetitions of encoding and variational layers. Each of the eight encoding layers encoded four features of the dataset, which created a lattice of 32 features across the four qubits. Finally, the measurement layer was included to produce two classical real-valued outputs. Fig. 8(a) shows this improvement in practice: the hybrid quantum neural network achieves a higher reward in only 8k steps, which remains inaccessible to a classical network of the same complexity even after 110k environment steps. Notably, these two networks were trained five times with varying initialization points, and the plot in Fig. 8(a) only shows the best runs of the hybrid and the classical models. The reader might notice that the hybrid model starts at a higher mean reward than the classical, a behavior observed elsewhere in the literature [13]. Additionally,

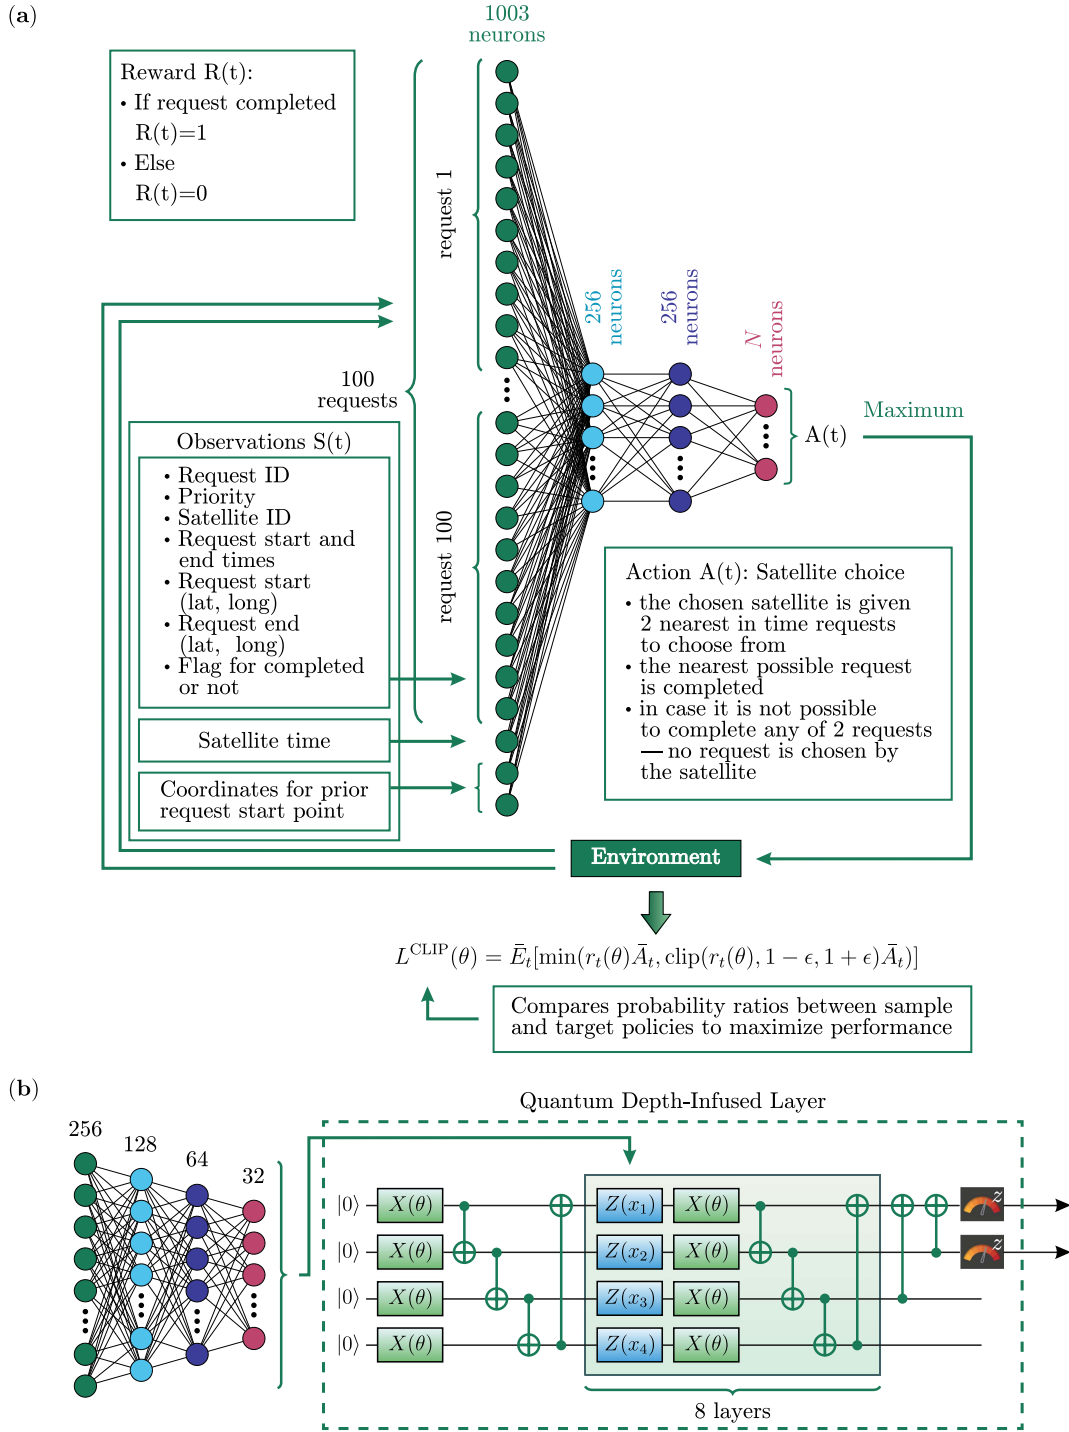


Fig. 5. (a) RL PPO model used to solve the mission planning problem. The state of the model encompasses the data of 100 requests, which are fed into the MLP agent to output an action (one request that is selected to complete), which then feeds into the environment (clip equation), generates an appropriate reward, and updates the state. (b) Quantum-hybrid RL model. A quantum circuit (left) is added to the beginning of the MLP Agent in the RL model (right) to incorporate quantum computation into the neural network.

the quick ascent aligns with the findings of [33], suggesting that quantum models can generalize from a few data points. Finally, Fig. 6 visualizes the solution chaining performance of the PPO algorithm.

4) *Hybrid AlphaZero*: The optimal results are expected from picking the best candidate from part of the solution pipeline. Specifically, in Section III-D, it was shown that RL is a powerful

algorithm that can be boosted in training and solution optimality through hybridization with quantum models. This section shows hybridized AlphaZero [16].

The classical AlphaZero was a RL algorithm developed by DeepMind that showed promise in solving difficult RL problems, such as chess, shogi, and go [16]. AlphaZero uses a computational tree of environment states whose values and

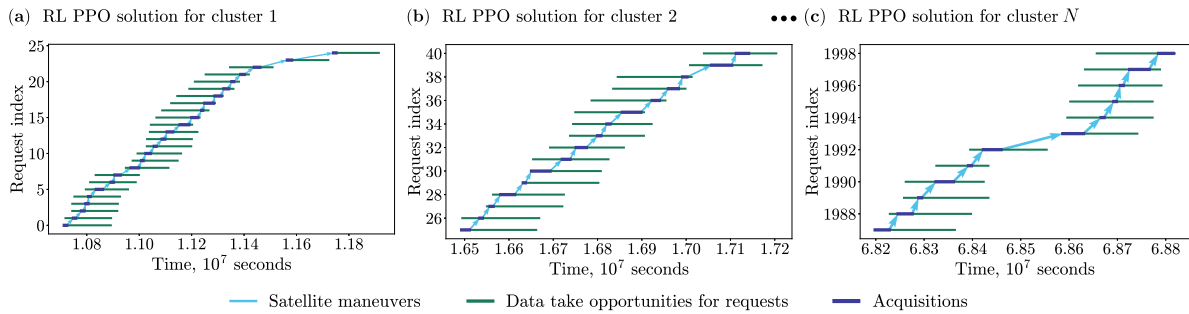


Fig. 6. Visualization of the chaining process used to complete requests within each cluster of the PPO algorithm.

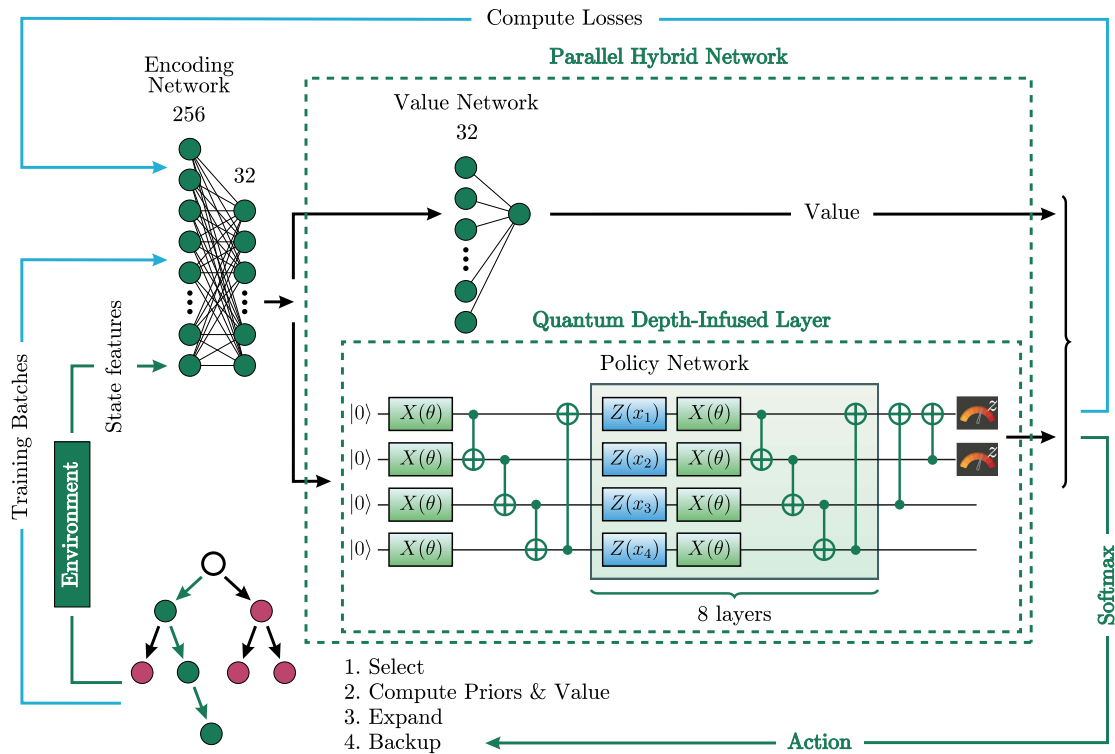


Fig. 7. Hybrid AlphaZero architecture used to achieve a near-optimal solution to the satellite mission planning. This architecture was motivated by the parallel hybrid networks (PHN) in [36] where information processing in quantum and classical models flow in parallel. The particular data encoding was based on the quantum depth-infused layer shown in [14].

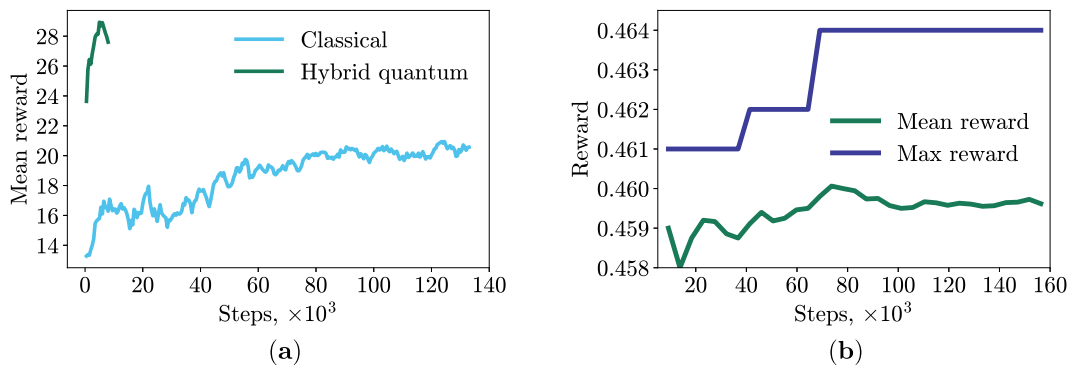


Fig. 8. (a) Rewards the classical and hybrid PPO RL agents achieve. The hybrid model reaches a higher reward at a significantly improved learning speed. (b) Agent reaches a maximum normalized reward of 0.464 after 70 000 steps of the simulation. This is compared to a total possible reward of 0.471, translating to an accuracy of 98.5%.

TABLE III
SUMMARY OF THE PERFORMANCE OF THE ALGORITHMS ON THE 2000 REQUESTS TASK PERFORMED WITH 2 SATELLITES

Model	Greedy only	Greedy + KMeans	RL (Hybrid AlphaZero)	Optimization
π_1 Completion %	75.8%	63.6%	98.5%	98.1%

The comparison demonstrates the greedy algorithm as a baseline, the hybrid alphazero, and integer optimization models.

probabilities are determined from the outputs of value and policy networks. The AlphaZero implementation in this work comprises four parts: a Monte Carlo tree search (MCTS), an encoding network, a policy network, and a value network.

The foundation of this model is akin to a MCTS model, which is used primarily for path prediction problems and board games based on strategy [34], [35]. The four major components of MCTS are selection, expansion, simulation, and backpropagation. The model begins at the tree's root and selects optimal child nodes until it reaches a leaf. Once it reaches the leaf, it expands the tree, creating a new child node. Then, the model simulates the remainder of the path-finding process from the newly created node and finally backpropagates with the newfound information to update the hyperparameters of the tree. The hybrid AlphaZero model uses MCTS with a parametrized quantum circuit as its policy network. The MLP agent of the model first recommends a state action. Then the action is applied to the environment, generating a reward and an updated state. However, in this case, the environment is the MCTS model; the recommended action is simulated in the model. Then, the backpropagation step of the MCTS is used to update the weights within the tree itself to refine its recommendations.

As each iteration of this loop occurs, there is also an outer loop in this algorithm. That outer loop is primarily concerned with the loss values being optimized; as a result, once the agent recommends an action, the outer loop takes training examples of a certain size and simulates the results over these batches to calculate the policy loss, value loss, and total loss. Once these values are calculated, the neural network weights are adjusted and fed back into the original loop, where the cycle repeats until the loss values are optimized.

Fig. 7 illustrates the inner workings of the hybrid AlphaZero model employed in this article. The state variables are passed into an initial encoding network with two hidden layers of sizes 256 and 32 neurons. The information is then passed in parallel [36] to 1) a single neuron using a fully connected layer and 2) a policy network as a PQC. The PQC resembles the one explained and implemented in the PPO model in Section III-D3. Fig 8(b) shows the training performance of this model (best out of five tries), which achieves a completion rate of $\pi_1 = 98.5\%$. Presenting the highest π_1 completion rate on the 2000 requests and 2 satellites dataset of any other model explored in this article. The completion rates of the highest performing algorithms from each section on this dataset are displayed in Table III.

IV. DISCUSSION

This work provided two classes of solutions to the scheduling problem of satellite mission planning: optimization and hybridized RL. From each class, the best-performing candidates were the integer optimization model and the hybrid AlphaZero, respectively. The dataset with 2 satellites and 2000 was used as a performance benchmark for the algorithms presented in

this work. The optimization and hybrid AlphaZero algorithms achieved π_1 completion rates of 98.1% and 98.5%, respectively, while the greedy algorithm only exhibited a 78.5% (63.6% with k -means clustering) completion rate. In the single satellite model, the optimization algorithm reached 100% completion on π_1 , π_2 , and π_3 requests and 96.2% completion on π_4 requests in 6 min. This work showed that by using RL and optimization models, it is possible to improve the results of mission planning that are otherwise obtained through simple greedy models. This work presents a step toward creating quantum-enhanced solutions in the space industry.

APPENDIX

A. Bunching Algorithms

Below are a few bunching algorithms used to cluster the data purely by sorting with respect to certain dataset features.

1) *DTO Bunching*: This algorithm first sorts the data by the DTO start times. After that, the data points are clustered by DTO overlap. In other words, all entries in any given cluster share at least one portion of their DTO windows with *all* other requests in the same cluster. As DTO windows are defined by the time when the satellite is in a position to capture the request, this method also functions as a form of geographical clustering due to the correlation of the location and DTO times.

2) *Priority Bunching*: In this algorithm, the requests are simply clustered by their priority ranks (1–4) and then ordered by the DTO start times. The utility of this lies greatly in the importance of the priority rankings, as the algorithm operates on the premise of each priority rank request being worth n times a request of the next highest priority rank, where n is the order of magnitude between two priority sets (for instance, a priority 1 request would be n^2 times as valuable as a priority 3 request).

3) *Bunch Sort*: This algorithm puts together the two clustering algorithms above, in a sense. First, the data are sorted by DTO windows, after which it is clustered by the DTO overlap. However, after that, the data are sorted within each cluster, this time by the end of their DTO windows and then by priority index. Consequently, data points with the same priority in the same cluster are ranked by the end of their DTO windows, maintaining the DTO structure within the priority indexing structure.

B. QUBO Formulation

Adiabatic quantum computers can approximately solve NP-hard problems, such as quadratic unconstrained binary optimization, faster than classical computers. Since many machine learning problems are also NP-hard, adiabatic quantum computers might be instrumental in training machine learning models efficiently in the post-Moore's law era. To solve problems on adiabatic quantum computers, they must be formulated as QUBO problems, which is possible by several techniques. This article formulated the problem as a QUBO problem, making

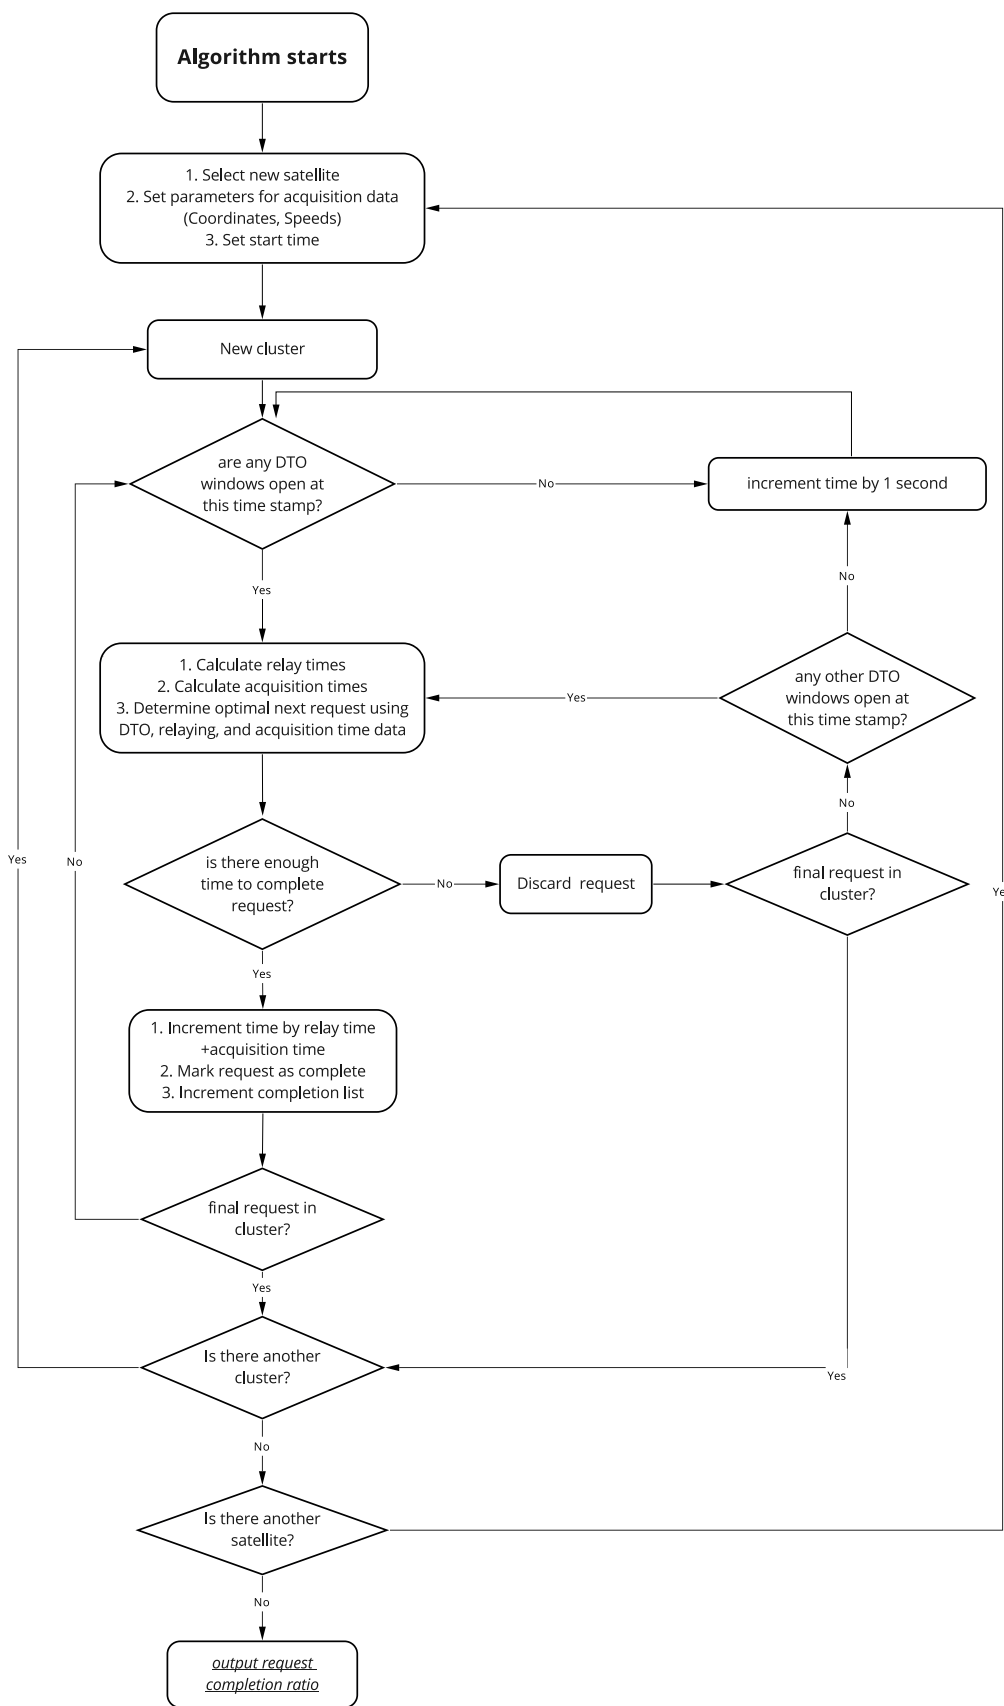


Fig. 9. Flowchart used to implement the greedy algorithm in Section III-B. The time-incrementing block references the “open DTO windows” query. Without accessible DTOs within a cluster, satellites remain idle under the greedy algorithm, with the simulation advancing until availability arises. Upon addressing all requests, satellites transition to the subsequent cluster and replicate the procedure. Note that the connection between one-second time increments and new clusters is implicit rather than explicit.

it conducive to being trained on adiabatic quantum computers. Since all the constraints will go into the cost function with some coefficients, they must be algebraically reformulated. Below are examples of how these constraints change according to the QUBO formulation

$$\forall f \in \mathbf{F}, \sum_{q \in \mathbf{Q}} x_{fq} \leq 1 \Rightarrow \left(\sum_{q \in \mathbf{Q}} x_{fq} - \frac{1}{2} \right)^2 \quad (21)$$

$$\forall q \in \mathbf{Q}, \sum_{f \in \mathbf{F}} x_{fq} \leq 1 \Rightarrow \left(\sum_{f \in \mathbf{F}} x_{fq} - \frac{1}{2} \right)^2 \quad (22)$$

$$\forall f \in \mathbf{F}, \sum_{\alpha \in \mathfrak{A}_f^{\text{start}}} y_{f\alpha} = \sum_{q \in \mathbf{Q}} x_{fq} \Rightarrow \quad (23)$$

$$\left(\sum_{\alpha \in \mathfrak{A}_f^{\text{start}}} y_{f\alpha} - \sum_{q \in \mathbf{Q}} x_{fq} \right)^2 \quad (24)$$

$$\forall q > 0, \sum_{f_1 \in \mathbf{F}} x_{f_1 q-1} \geq \sum_{f_2 \in \mathbf{F}} x_{f_2 q} \Rightarrow \quad (25)$$

$$\sum_{f \in \mathbf{F}} (x_{f_2 q} - x_{f_1 q-1}) + \sum_i 2^i z_i = 0$$

where z_i is a slack binary variable, and the number of slack variables depends on the problem. Thus, the idea of the QUBO formulation is to reduce all constraints to linear equalities, square them, and then pose as penalties in the cost function. The more rigorous analysis of the slack variable implementation is provided in [37]. QUBO formulation is written in general terms as

$$Q = Q_0 + \beta C, \quad C \geq 0 \quad (26)$$

where Q is a cost function, Q_0 is the objective function of the initial problem, β are some coefficients that must be picked up, and C are nonnegative definite quadratic constraints. The solution then takes the following form:

$$Q[q] = \min_{\kappa} Q[\kappa]. \quad (27)$$

In the case where $C[q] > 0$, it is possible to choose such β , that if q is unconstrained optimum, then it must be feasible. Beta is chosen such that if q is not feasible, then

$$Q[q] > Q[\text{feas}] = Q[0] = 0 \quad (28)$$

where $Q[\text{feas}]$ means the initial zero solution.

The initial objective function that was used for the classical solution will take the following form for the quantum solution:

$$Q_0 = - \sum_{f \in \mathbf{F}} \left(\sum_{q \in \mathbf{Q}} J_f x_{fq}^2 - \sum_{\alpha \in \mathfrak{A}_f^{\text{start}}} \gamma_{f\alpha} y_{f\alpha}^2 \right) \rightarrow \min. \quad (29)$$

Also, the odds β are chosen so that there is a better solution when the constraints are violated

$$C[q] > 0 \Rightarrow Q[q] \geq Q_0[q] + \beta > 0. \quad (30)$$

Accordingly, in this case it is necessary to take $\beta > -Q_0[q]$. The boundaries of q must be estimated, as the precise values are undetermined. The worst estimate was chosen as the following:

$$\beta > -Q[q] = - \left(\min_{\kappa} Q_0[\kappa] \right) = - \left(- \sum_f J_f \right) = \sum_f J_f. \quad (31)$$

C. Greedy Algorithm

Fig. 9 showcases the logic of the greedy algorithm from Section III-B.

REFERENCES

- [1] P. Chini, G. Giambene, and S. Kota, "A survey on mobile satellite systems," *Int. J. Satell. Commun. Netw.*, vol. 28, pp. 29–57, 2009.
- [2] E. Dziadczyk, W. Zabierowski, and A. Napieralski, "Satellite navigation system gps," in *Proc. 9th Int. Conf. – Experience Design. Appl. CAD Syst. Microelectron.*, 2007, pp. 504–506.
- [3] B. Jaksic, D. Miljkovic, V. Maksimovic, M. Petrovic, and B. Gvozdic, "Satellite television transmission in the world - broadcasting systems and standards," *Acta Scientiarum. Technol.*, vol. 42, no. 1, 2020, Art. no. e44957.
- [4] R. Grasset-Bourdel, G. Verfaillie, and A. Flipo, "Planning and replanning for a constellation of agile earth observation satellites," in *ICAPS-11 Workshop Scheduling Plan. Appl.*, Freiburg, Germany, 2011.
- [5] Y. Huang, Z. Mu, S. Wu, B. Cui, and Y. Duan, "Revising the observation satellite scheduling problem based on deep reinforcement learning," *Remote Sens.*, vol. 13, no. 12, 2021, Art. no. 2377.
- [6] M. Lemaitre, G. Verfaillie, F. Jouhaud, J.-M. Lachiver, and N. Bataille, "Selecting and scheduling observations of agile satellites," *Aerosp. Sci. Technol.*, vol. 6, pp. 367–381, 2002.
- [7] H. A. Khojah and M. A. Mosa, "Multi-objective optimization for multi-satellite scheduling task," *J. Soft Comput. Exploration*, vol. 3, pp. 19–30, 2022.
- [8] J. Biamonte, P. Wittek, N. Pancotti, P. Rebentrost, N. Wiebe, and S. Lloyd, "Quantum machine learning," *Nature*, vol. 549, no. 7671, 2017, Art. no. 195.
- [9] V. Dunjko and H. J. Briegel, "Machine learning & artificial intelligence in the quantum domain: A review of recent progress," *Rep. Prog. Phys.*, vol. 81, no. 7, 2018, Art. no. 074001.
- [10] A. Melnikov, M. Kordzanganeh, A. Alodjants, and R.-K. Lee, "Quantum machine learning: From physics to software engineering," *Adv. Phys.: X*, vol. 8, no. 1, 2023, Art. no. 2165452.
- [11] G. Wu, H. Wang, W. Pedrycz, H. Li, and L. Wang, "Satellite observation scheduling with a novel adaptive simulated annealing algorithm and a dynamic task clustering strategy," *Comput. Ind. Eng.*, vol. 113, pp. 576–588, 2017.
- [12] T. Stollenwerk, V. Michaud, E. Lobe, M. Picard, A. Basermann, and T. Botter, "Image acquisition planning for Earth observation satellites with a quantum annealer," Jul. 10, 2023, *arXiv:2006.09724*.
- [13] A. Sagingalieva et al., "Hyperparameter optimization of hybrid quantum neural networks for car classification," Jul. 10, 2023, *arXiv:2205.04878*.
- [14] A. Sagingalieva, M. Kordzanganeh, N. Kenbayev, D. Kosichkina, T. Tomashuk, and A. Melnikov, "Hybrid quantum neural network for drug response prediction," *Cancers*, vol. 15, no. 10, 2023, Art. no. 2705.
- [15] M. Perelshtein et al., "Practical application-specific advantage through hybrid quantum computing," Jul. 10, 2023, *arXiv:2205.04858*.
- [16] D. Silver et al., "A general reinforcement learning algorithm that masters chess, shogi, and go through self-play," *Sci.*, vol. 362, no. 6419, pp. 1140–1144, 2018.
- [17] M. Kordzanganeh et al., "Benchmarking simulated and physical quantum processing units using quantum and hybrid algorithms," *Adv. Quantum Technol.*, vol. 6, 2023, Art. no. 2300043.
- [18] X. Jin and J. Han, *K-Means Clustering*. Berlin, Germany: Springer US, 2010, pp. 563–564.
- [19] J. MacQueen, "Classification and analysis of multivariate observations," in *Proc. 5th Berkeley Symp. Math. Statist. Probability*, 1967, pp. 281–297.

- [20] S. Kilic, "An optimization approach for the daily photograph selection of earth observation satellites," *Havacılık ve Uzay Teknolojileri Dergisi*, vol. 12, 2019.
- [21] A. Land and A. Doig, "An automatic method of solving discrete programming problems," *Econometrica: J. Econometric Soc.*, pp. 497–520, 1960.
- [22] L. Maisonobe, V. Pommier, and P. Parraud, "Orekit: An open source library for operational flight dynamics applications," in *Proc. 4th Int. Conf. Astrodynamics tools Techn.* ESAC Madrid, 2010, pp. 3–6.
- [23] Gurobi Optimization, LLC, "Gurobi optimizer reference manual," 2022.
- [24] E. Tosatti, "Optimization using quantum mechanics: Quantum annealing through adiabatic evolution," *J. Phys. A: Math. Theor.*, vol. 41, 2008, Art. no. 209801.
- [25] M. W. Johnson et al., "Quantum annealing with manufactured spins," *Nature*, vol. 473, no. 7346, pp. 194–198, 2011.
- [26] J. Schulman, F. Wolski, P. Dhariwal, A. Radford, and O. Klimov, "Proximal policy optimization algorithms," Jul. 10, 2023, *arXiv:1707.06347*.
- [27] S. Jerbi, C. Gyurik, S. Marshall, H. Briegel, and V. Dunjko, "Parameterized quantum policies for reinforcement learning," *Adv. Neural Inf. Process. Syst.*, vol. 34, pp. 28362–28375, 2021.
- [28] A. Skolik, S. Jerbi, and V. Dunjko, "Quantum agents in the Gym: A variational quantum algorithm for deep Q-learning," *Quantum*, vol. 6, 2022, Art. no. 720.
- [29] A. Senokosov, A. Sedykh, A. Sagingalieva, and A. Melnikov, "Quantum machine learning for image classification," Jul. 10, 2023, *arXiv:2304.09224*.
- [30] A. Sedykh et al., "Quantum physics-informed neural networks for simulating computational fluid dynamics in complex shapes," Jul. 10, 2023, *arXiv:2304.11247*.
- [31] M. Benedetti, E. Lloyd, S. Sack, and M. Fiorentini, "Parameterized quantum circuits as machine learning models," *Quantum Sci. Technol.*, vol. 4, no. 4, 2019, Art. no. 043001.
- [32] M. Schuld and N. Killoran, "Quantum machine learning in feature Hilbert spaces," *Phys. Rev. Lett.*, vol. 122, no. 4, 2019, Art. no. 040504.
- [33] M. C. Caro et al., "Generalization in quantum machine learning from few training data," *Nature Commun.*, vol. 13, no. 1, 2022, Art. no. 4919.
- [34] G. Chaslot, S. Bakkes, I. Szita, and P. Spronck, "Monte-Carlo tree search: A new framework for game AI," in *Proc. AAAI Conf. Artif. Intell. Interactive Digit. Entertainment*, vol. 4, pp. 216–217, 2008.
- [35] B.-D. Lee, "Monte-Carlo tree search applied to the game of tic-tac-toe," *J. Korea Game Soc.*, vol. 14, pp. 47–54, 2014.
- [36] M. Kordzanganeh, D. Kosichkina, and A. Melnikov, "Parallel hybrid networks: An interplay between quantum and classical neural networks," Jul. 10, 2023, *arXiv:2303.03227*.
- [37] A. Pakhomchik, S. Yudin, M. Perelshtein, A. Alekseyenko, and S. Yarkoni, "Solving workflow scheduling problems with QUBO modeling," Jul. 10, 2023, *arXiv:2205.04844*.



Serge Rainjonneau received the Ph.D. degree in simulation, AI & image synthesis from the Paul Sabatier University of Toulouse, Toulouse, France, in 1992, focusing on the implementation of behavioral simulations in computer graphics.

He assumed a Research and Development Engineer role with Thales Alenia Space, France, from 1992 to 1996 on different aspects of satellite mission performances. From 1996 to 2014, he was responsible for satellite constellation engineering and mission performances for different Telecom and Earth observation programs still at Thales Alenia Space, France. From 2014 up to 2018, as a Mission Expert, he represented Thales Alenia Space at the French Institute of Technology (IRT Saint-Exupery) leading activities on both satellite on-board autonomy and on-ground mission planning in reactive loop. From 2018 up to 2020, he was the Head of Mission Pole of Thales Alenia Space for Earth Observation Division. Since the beginning of 2021, he leads the Thales Alenia Space product line architecture dedicated to Earth Observation ground segments. In the frame of these activities, he leads several studies involving AI and Quantum Computing for difficult problems on combinatorial optimization like mission plan generation for constellations of agile Earth observation satellites.



Igor Tokarev received the Graduate degree in computer security in 2020. Since 2022, he has been working with Terra Quantum, applying quantum machine learning algorithms to real-world problems. His research interest include deep learning, reinforcement learning, quantum computing, and cybersecurity.



Sergei Iudin received the bachelor's degree specialized on "Solving workflow scheduling problems with QUBO modeling" at the Department of General and Applied Physics in 2022.

Since 2021, he has been exploring optimization algorithms for noisy intermediate-scale quantum devices, including gate-based approaches and quantum annealing. He currently works as a Python Developer and optimization specialist at Terra Quantum AG, St. Gallen, Switzerland, involved in research in classical and quantum optimization.



quantum algorithms.

Saaketh Rayaprolu received the master of science degree in quantum information science from the University of Southern California, Los Angeles, CA, USA, in 2023.

Prior to graduating, he was employed as a Quantum Machine Learning intern at Terra Quantum in 2022, working to apply quantum solutions to complex computational problems, and was one of the victors of the quantum computing hackathon hosted at the International Centre for Theoretical Physics in 2023. His research interests include mathematics, data science,




at Hubble HQ, deep tech commercialization at Sia Partners, and AI/DevOps software development at Accenture.

Currently serving as the Director of Growth at Terra Quantum, Karan leverages his expertise to develop high-value applications and spearhead pioneering quantum technology innovations aimed at addressing critical global challenges. His entrepreneurial endeavors encompass various consumer goods and services-based businesses. Notably, he co-created "Entangle Climate" and co-founded "Aedis," initiatives that harness technology to mitigate climate change and explore emerging domains.



Daria Lemtiuzhnikova works as a Senior Researcher with Terra Quantum. Her Ph.D. thesis in theoretical computer science and mathematical cybernetics was on "Dimension reduction for large problems with sparse matrices" in 2018.



Miras Koblan is a medalist of national physics competitions and international olympiads as a member of Kazakhstan national team. Since 2022, he has been working in Terra Quantum as a junior machine learning developer, specializing in reinforcement learning and quantum machine learning algorithms.



Egor Barashov received the master's degree (diploma with Hons.) in 2023 in the Department of Physical and Mathematical Control Methods.

He began his scientific activity by working in the laboratory of scheduling theory and discrete optimization. Since 2022, he has worked with classical and quantum methods for solving optimization problems at Terra Quantum AG, St. Gallen, Switzerland.



Mo Kordzanganeh received the first Class Master of Physics degree in 2021 from the school of Astronomy and Astrophysics at The University of Manchester, Manchester, U.K., where he developed quantum machine learning techniques for radio astronomy under the supervision of Prof. Anna Scaife. The results of this endeavor were published in the Machine Learning for Physical Sciences workshop at the 2021 NeurIPS conference.

He is a Quantum Machine Learning Scientist with Terra Quantum AG, where he has been working at the intersection of quantum physics and machine learning, trying to resolve the pain-points of classical machine learning with quantum computing.



Markus Pflictsch is the CEO, Chairman, and Founder of Terra Quantum. He is a dedicated quantum physicist, senior financial executive, and deep tech entrepreneur. Since he was 14 years old, he has been a quantum physics enthusiast. He built on that passion, studying mathematics and physics at RWTH Aachen, and researched in the area of quantum field theory at CERN. His curiosity drew him to business and finance, and he subsequently joined Boston Consulting Group. Later on, he assumed management responsibilities as a Senior Executive in various corporates

including Deutsche Bank and UniCredit, as well as positions as Chief Financial Officer in owner-managed digital and high-tech companies. Ultimately, he decided to leverage his expertise in quantum physics, business, and finance to become an entrepreneur in the quantum technologies space. In 2018, he founded Terra Quantum AG, a deep tech pioneer commercializing quantum tech applications. Markus is a senior advisor to private equity funds and a member of the "Baden-Badener Unternehmensgespräche (BBUG)."



Alexey Melnikov received the Ph.D. degree in physics from the University of Innsbruck, Innsbruck, Austria, in 2018, focusing on the intersection of physics and machine learning, including quantum machine learning.

Post-Ph.D., he assumed postdoctoral roles with the University of Innsbruck and the University of Basel, Switzerland, from 2018 to 2020, exploring the integration of AI in physics. As a Visiting Assistant Professor, he served the National Tsing Hua University, Hsinchu, Taiwan. In 2019, he and his coauthors

were honored with the prestigious Cozzarelli Prize from the National Academy of Sciences, USA, in the Physical and Mathematical Sciences category. Since 2020, he has led the Machine Learning at Terra Quantum AG, St. Gallen, Switzerland, directing a research team that leverages quantum machine learning for industrial benefit.