



VQC-based reinforcement learning with data re-uploading: performance and trainability

Rodrigo Coelho^{1,2} · André Sequeira^{1,2,3} · Luís Paulo Santos^{1,2,3}

Received: 28 January 2024 / Accepted: 9 August 2024 / Published online: 28 August 2024
© The Author(s), under exclusive licence to Springer Nature Switzerland AG 2024

Abstract

Reinforcement learning (RL) consists of designing agents that make intelligent decisions without human supervision. When used alongside function approximators such as Neural Networks (NNs), RL is capable of solving extremely complex problems. Deep Q-Learning, a RL algorithm that uses Deep NNs, has been shown to achieve super-human performance in game-related tasks. Nonetheless, it is also possible to use Variational Quantum Circuits (VQCs) as function approximators in RL algorithms. This work empirically studies the performance and trainability of such VQC-based Deep Q-Learning models in classic control benchmark environments. More specifically, we research how data re-uploading affects both these metrics. We show that the magnitude and the variance of the model's gradients remain substantial throughout training even as the number of qubits increases. In fact, both increase considerably in the training's early stages, when the agent needs to learn the most. They decrease later in the training, when the agent should have done most of the learning and started converging to a policy. Thus, even if the probability of being initialized in a Barren Plateau increases exponentially with system size for Hardware-Efficient ansatzes, these results indicate that the VQC-based Deep Q-Learning models may still be able to find large gradients throughout training, allowing for learning.

Keywords Reinforcement learning · Quantum computing · Neural Networks · Variational Quantum Circuits · Quantum Machine Learning

1 Introduction

1.1 Context

With quantum computers still in the *Noisy Intermediate Scale Quantum (NISQ)* era, algorithms with theoretical complexity advantages over the best-known classical algorithms, such as Shor's algorithm to factor large integers (Shor 1999), cannot

be implemented to solve practical problems (Preskill 2018). Thus, much of nowadays quantum computing research focuses on finding near-term quantum advantages. A field of quantum computing that has attracted considerable interest as one of the most promising applications for achieving short-term quantum advantage is *Quantum Machine Learning (QML)*. This is mainly due to *Variational Quantum Circuits (VQCs)* (Cerezo et al. 2021).

VQCs are quantum circuits that depend on free parameters, which are iteratively updated by a classical optimizer to minimize an objective function estimated from measurements (Ostaszewski et al. 2021). VQCs are especially suitable for NISQ devices, given their lower demands in the number of qubits and circuit depth, which mitigates the effect of noise (Cerezo et al. 2021). Consequently, integrating VQCs into machine learning algorithms has become an intriguing field of study. VQCs used in such a context have been extensively researched in both supervised (Mitarai et al. 2018; Schuld and Killoran 2019; Schuld et al. 2020; Farhi and Neven 2018) and unsupervised (Coyle et al. 2020; Zoufal et al. 2021) QML.

✉ Rodrigo Coelho
rodrigocoelho140@gmail.com

André Sequeira
andresequeira410@gmail.com

Luís Paulo Santos
psantos@di.uminho.pt

¹ Department of Informatics, University of Minho, Braga, Portugal

² HASLab, INESC TEC, Braga, Portugal

³ International Nanotechnology Laboratory (INL), Braga, Portugal

A RL agent learns from interaction with its surrounding environment in a feedback loop, relying solely on a reward signal to guide its actions. The concept of learning from interaction forms the bedrock of nearly all theories of learning and intelligence (Sutton and Barto 2018), with some even conjecturing that rewards alone can pave the path to achieving general AI agents (Silver et al. 2021). RL techniques, with the advent of deep learning (LeCun et al. (2015)), consecrated a major impact in tasks like Chess (Silver et al. 2017), Go (Silver et al. 2016) and video-games (Mnih et al. 2015), matching or exceeding the performance levels of humans. It has also been applied to practical real-world problems, such as autonomous driving (Sallab et al. 2017), robotics (Kober et al. 2013) and finance (Hambly et al. 2023).

Just like VQCs have found applications as function approximators in both supervised and unsupervised ML domains, they can be effectively utilized in the same capacity in RL scenarios, serving as alternatives to DNNs. These quantum variational approaches in the context of RL have started gaining traction only recently, as evidenced by several recent studies (Chen et al. 2020; Lockwood and Si 2020; Skolik et al. 2022; Jerbi et al. 2021; Sequeira et al. 2022; Skolik et al. 2023), which are further discussed on the next subsection.

1.2 Related work

It was recently shown that the data encoding strategy used in the quantum circuit influences the expressive power of the quantum model (Schuld et al. 2021). Specifically, the authors demonstrate that a VQC may be written as a Partial Fourier Series, where the accessible data frequencies are determined by the eigenvalues of the data encoding Hamiltonians. Furthermore, quantum models can access increasingly rich frequency spectra by repeating simple data encoding gates multiple times either in series or in parallel, a technique named *Data Re-uploading* (Pérez-Salinas et al. 2020).

VQCs have been used as function approximators in two different RL algorithms: Deep Q-learning (Chen et al. 2020; Lockwood and Si 2020; Skolik et al. 2022) and Reinforce with baseline (Jerbi et al. 2021; Sequeira et al. 2022). For the VQC-based Reinforce algorithms, Jerbi et al. (2021) employed a VQC with data re-uploading and assessed the performance of models in benchmark environments. Conversely, Sequeira et al. (2022) used a simpler, hardware-efficient ansatz without data re-uploading, comparing its performance and trainability against classical DNNs. Intriguingly, the impact of data re-uploading on model trainability remains unexplored. Given that data re-uploading increases the VQC's circuit depth and expressivity, it could have negative effects on the trainability, as suggested by McClean et al. (2018) and Holmes et al. (2022).

The picture is even less clear when it comes to VQC-based Deep Q-Learning algorithms. Both Chen et al. (2020) and Lockwood and Si (2020) used simple VQCs without data re-uploading, centering their studies on performance in benchmark environments and comparisons with DNNs. On the other hand, Skolik et al. (2022) integrated data re-uploading and analyzed its impact on performance. Yet, to the best of our knowledge, no research has been done on the trainability of VQC-based Deep Q-Learning models.

This paper aims to directly address these questions. Specifically, we seek to investigate the impact of data re-uploading on both the performance and trainability of VQC-based Deep Q-Learning models. Our approach involves empirically analyzing the average return of two different hardware-efficient VQCs in benchmark environments and the behavior of their gradient's norm and variance. Additionally, we also verify how these metrics are affected by an increasing system size.

1.3 Contributions

The contributions of this paper are organized as follows:

- Empirically verified the importance of data re-uploading and trainable output and input scaling in enhancing the performance of VQC-based Deep Q-learning models in OpenAI Gym's CartPole-v0 and Acrobot-v1.
- Achieved considerable performance in the Acrobot-v1 environment, which was previously untapped for VQC-based Deep Q-learning models.
- Empirically verified that the gradient's variance of the quantum models increases in the early stages of training in the aforementioned environments. Furthermore, we verified that the gradient's variance increases when data re-uploading is used.
- Empirically identified a correlation between the moving targets inherent to Deep Q-learning (Mnih et al. 2015), gradient magnitude and variance, and performance.

1.4 Outline

The remainder of the manuscript is organized as follows. Section 2 introduces the mathematical framework behind most RL problems, the Q-Learning algorithm and the result of its combination with Deep Neural Networks. Section 3 then introduces the main building blocks of Variational Quantum Circuits and how these can be used as function approximators in Deep Q-Learning. Section 4 proposes the methodologies for analyzing performance and trainability and applies them to VQC-based Deep Q-Learning models in benchmark environments. Finally, Section 5 finishes the paper with some final remarks and suggestions for future work.

2 Reinforcement learning

2.1 Introduction

Any RL problem has two primary constituents: the Agent and the Environment. The agent is the decision-maker and the representation of that which learns. The environment is everything that surrounds the agent and is affected by its decisions. The agent must learn which actions to take by interacting with the environment and receiving a reward that indicates the goodness of that particular action concerning the goal, see Fig. 1.

The mathematical framework behind this interface and most RL problems is known as a *Markov Decision Process (MDP)* (Sutton and Barto 2018), which is expressed as a tuple $\langle \mathcal{S}, \mathcal{A}, \mathcal{P}, R \rangle$ (Sutton and Barto 2018), encompassing a finite set of states \mathcal{S} , actions \mathcal{A} , a state transition probability matrix \mathcal{P} , and the reward function R .

A crucial metric in RL is the *Return*, that is, the cumulative discounted reward starting from time step t , such that:

$$G_t = R_{t+1} + \gamma R_{t+2} + \gamma^2 R_{t+3} + \dots = \sum_{k=0}^{\infty} \gamma^k R_{t+k+1} \quad (1)$$

where $\gamma \in [0, 1]$ is the *discount factor*, which determines how much long-term rewards are valued over short-term




Fig. 1 The Agent-Environment Interface: The agent interacts with the environment at time step t by taking action A_t . The environment then changes to state S_{t+1} and produces reward R_{t+1} , which are both passed back to the agent so that it can decide the next action. The dotted lines indicate that this process repeats itself. Inspired by Sutton and Barto (2018)

rewards. The behavior of the agent is given by the *policy* $\pi(a|s)$, which is a probability distribution over actions given states. It determines the probability the agent will take each action a from a particular state s . The objective of a RL agent becomes clear: to identify the policy that maximizes the expected return $\max_{\pi} \mathbb{E}_{\pi}[G_t]$. This policy is typically referred to as the *optimal policy*. It is known that, for any MDP, there exists at least one deterministic optimal policy π_* that is better than or equal to all other policies, $\pi_* \geq \pi, \forall \pi$ (Sutton and Barto 2018). Consequently, the goal is for the policy of the agent to converge to this deterministic optimal policy as training goes on.

There are several methods for finding a (near-) optimal policy. We first distinguish between *model-based* and *model-free* methods (Sutton and Barto 2018; Kaelbling et al. 1996). In model-based methods, the agent has access to (or learns an approximated version of) the model and can use it to plan. In model-free methods, the agent doesn't have access to a model. These methods are simpler and easier to implement, but also usually suffer from a higher sample complexity, since the agent has to deal with the extra burden of not knowing the dynamics of the environment.

Then, within the model-free methods, we need to distinguish between *value-based* and *policy-based* methods. The former use value functions as guides in the search for optimal policies (Sutton and Barto 2018), such as Q-learning (Watkins and Dayan 1992). Policy-based methods learn the policy directly. They typically work by parameterizing the policy $\pi_{\theta}(a|s)$, defining an objective function that serves as some performance measure $J(\pi_{\theta})$ and optimizing the parameters θ by performing gradient ascent on an approximation of the gradient $\nabla J(\pi_{\theta})$, which is obtained using Monte Carlo samples. Consequently, they are commonly referred to as *policy gradient* methods. The applicability of both value-based and policy-based methods vastly depends on the problem itself.

2.2 Q-learning

The *Action-Value Function* $q_{\pi}(s, a)$ defines how good taking a certain action from a particular state is. It is defined as:

$$q_{\pi}(s, a) = \mathbb{E}_{\pi}[G_t | S_t = s, A_t = a] \quad (2)$$

Furthermore, in any MDP, there exists an optimal action-value function $q_*(s, a)$, that corresponds to the maximum action-value function over all policies

$$q_*(s, a) = \max_{\pi} q_{\pi}(s, a) \quad (3)$$

This function is achieved by all the optimal policies (Sutton and Barto 2018). Actually, once the optimal action-value

function is known, the optimal policy can be found by maximizing over it:

$$\pi_*(a|s) = \begin{cases} 1 & \text{if } a = \arg \max_{a \in \mathcal{A}} q_*(s, a) \\ 0 & \text{otherwise} \end{cases} \quad (4)$$

Q-Learning is a model-free value-based algorithm that approximates the optimal action-value function and then, according to Eq. 4, a (near) optimal policy (Watkins and Dayan 1992).

Being model free, Q-Learning agents learn from experience — finite sequences of states, actions, and rewards, in the format $\{s, a, r, s'\}$, which are called *trajectories* (Sutton and Barto 2018), or *episodes* if the environment is episodic, which will be the case in this paper. To generate these episodes, some behavior policy has to be followed. Choosing this policy implies taking into consideration the *exploration-exploitation tradeoff*. On the one hand, the agent has to explore the state space to find high-value states. On the other hand, the agent has to exploit the acquired knowledge, taking actions that it knows yield high rewards and lead to good states. As a result, typically, an ϵ -greedy policy, which chooses a random action with probability ϵ and the greedy action with probability $1 - \epsilon$, is used (Sutton and Barto 2018). It is usual to decay ϵ over time.

To approximate the optimal Q -values, Q-Learning uses a method known as *Temporal-Difference Learning (TD-Learning)* — it learns from raw experience but updates estimates with other learning estimates without waiting for an outcome — *bootstrapping*. More specifically, the action-value of a certain state and a particular action is updated as such:

$$Q(S, A) \leftarrow Q(S, A) + \alpha [R + \gamma \max_a Q(S', a) - Q(S, A)] \quad (5)$$

2.3 Deep q-learning

The previous subsection introduces Q-Learning, which uses lookup tables of (at most) dimensions $\mathcal{S} \times \mathcal{A}$ to assign values to state-action pairs encountered by the agent — *tabular regime*. Consequently, the algorithm is inadequate for handling large and complex real-world problems, since these tables become impossible to implement with current technology due to memory and complexity constraints.

There are several possible solutions for this problem, such as using state-abstraction (Andre and Russell 2002) and/or dimensionality reduction techniques (Sorzano et al. 2014). This work focuses on the most popular solution — using *Function Approximation*. The idea is to generalize from seen states to unseen states by approximating a target function.

Consequently, a new function is defined:

$$\hat{q}(s, a, \mathbf{w}) \approx q_\pi(s, a) \quad (6)$$

where \mathbf{w} is a vector of parameters that must be updated using TD-Learning or MC. It is important to refer that function approximation is only useful if $\dim(\mathbf{w}) \ll |\mathcal{S} \times \mathcal{A}|$, otherwise it would be the same as the tabular regime.

There are many function approximators, each with pros and cons (Busoniu et al. 2017). Nonetheless, the most used function approximators are Neural Networks (NNs), particularly those with two or more hidden layers of neurons, known as Deep Neural Networks (DNNs), due to their capacity to approximate intricate and complex functions (LeCun et al. 2015). When DNNs are used as function approximators for the Q -function in a Q-Learning framework, the resultant algorithm is known as Deep Q-Learning or Deep Q-Networks (DQN) (Mnih et al. 2015). This algorithm introduces two novel improvements over previous approaches: *experience replay* and *target network*. The former consists of using a *Replay Buffer*, where all experiences — interactions with the environment in the format $(s_t, a_t, r_{t+1}, s_{t+1})$ — are stored. Then, batches of these experiences are randomly sampled to compute the parameter updates. Consequently, temporal correlations between the samples are broken, making the data appear more *Independent and Identically Distributed (IID)*, which stabilizes training. The *Target Network* consists of a second NN that is used to compute the TD-targets. The target network has “frozen” weights that are sporadically (every C steps) updated with a copy of the parameters of the online network. Thus, the targets are more stationary, which also stabilizes training. However, there is a tradeoff here between speed of convergence and stability which must be accounted for when choosing C .

3 Quantum variational deep q-learning

Quantum Variational Deep Q-learning, or VQC-based Deep Q-Learning, uses VQCs as function approximators in a Q-Learning framework, instead of the typically used DNNs. VQCs are quantum circuits that depend on free parameters, which are iteratively updated by a classical optimizer to minimize an objective function estimated from measurements (Ostaszewski et al. 2021). In the following three subsections, we will go over the main details of the VQCs used throughout this work.

3.1 Data encoding

Classical data has to be encoded into a quantum state so that it can be processed by a quantum computer. This paper uses *Angle Encoding*. According to this technique,

Algorithm 1 Deep Q-Learning (Mnih et al. 2015).

```

1: Algorithm parameters: step size  $\alpha \in [0, 1]$ , small  $\epsilon \in [0, 1]$ 
2: Initialize replay buffer  $\mathcal{D}$  to capacity  $N$ 
3: Initialize action-value function  $Q$  with random weights  $\theta$ 
4: Initialize target action-value function  $\hat{Q}$  with weights  $\theta^- = \theta$ 
5: for episode= 1,  $M$  do
6:   Initialize  $s_1$ 
7:   for  $t = 1, T$  do
8:     Choose  $a_t$  from  $s_t$  using  $\epsilon$ -greedy policy
9:     Take action  $a_t$ , observe  $r_{t+1}, s_{t+1}$ 
10:    Store experience  $(s_t, a_t, r_{t+1}, s_{t+1})$  in replay memory  $\mathcal{D}$ 
11:    Sample random mini-batch of transitions  $(s, a, r, s')$  from  $\mathcal{D}$ 
12:    Compute the TD-targets using target network (with old, fixed parameters  $\theta^-$ )
13:    Perform a gradient descent step on
      
$$\left[ r + \gamma \max_{a'} \hat{Q}(s', a'; \theta_i^-) - Q(s, a; \theta_i) \right]^2$$

      w.r.t parameters  $\theta$ 
14:    Every  $C$  steps update  $\hat{Q} = Q$ 
15:    Initialize  $s_t$ 
16:   end for
17: end for

```

each component of the input data x is encoded by a single qubit using arbitrary Pauli rotations R_X, R_Y, R_Z , where the angle is given by the component itself after some classical pre-processing, e.g., normalization. To allow for adaptive frequency matching between the function output by the VQC and the target function, each feature will be multiplied by a classical trainable weight, a technique known as trainable input scaling (Pérez-Salinas et al. 2020). Thus, given the input data $x = \{x_0, x_1, \dots, x_{n-1}\}$, the resulting quantum state is given by:

$$|x\rangle = \bigotimes_{i=0}^{n-1} R_\alpha(\phi(x_i \times \lambda_i))|0_i\rangle \tag{7}$$

where $R_\alpha \in \{R_X, R_Y, R_Z\}$, ϕ is some classical pre-processing function and $|0_i\rangle$ is the i^{th} qubit initialized in state $|0\rangle$.

This technique has pros and cons. On the one hand, it allows the encoding of a given input vector using circuits of depth 1 (Sequeira et al. 2022). On the other hand, the number of qubits grows linearly with the number of features of the input vector, which limits the dimensionality of the inputs one can encode due to the reduced number of qubits in NISQ devices.

3.2 Ansatz

A set of problem-agnostic ansatzes called *Hardware-Efficient Ansatzes* were considered in this work. These allow for implementations with reduced circuit depth by bringing correlated qubits together for depth-reduction (Cerezo et al. 2021). For this work, the ansatzes selected are those from Skolik et al. (2022) and the Universal Quantum Classifier

(UQC) (Pérez-Salinas et al. 2020), depicted in Fig. 2. For simplicity, the Skolik’s architecture will be referred to as *Skolik Data Re-Uploading* if data re-uploading is used and *Skolik Baseline* otherwise.

However, there is a tradeoff one needs to take into account when using hardware-efficient VQCs. Since these ansatzes take no inspiration from the structure of the problem itself, they need to be highly expressive so they can be applied to any generic task (Bilkis et al. 2021). However, this expressiveness comes at the cost of reduced trainability (Holmes et al. 2022). This problem, known as the *Barren Plateau Phenomenon*, is explained in Section 3.4.

3.3 Cost function and gradient-based learning

To use a VQC to solve a problem, one has to first encode the problem into a cost function, such that solving the problem corresponds to finding the global optima of said function. Let $f_\theta(x)$ be the expectation value of some observable \hat{O} as follows:

$$f_\theta(x) = \langle 0|U^\dagger(x, \theta)\hat{O}U(x, \theta)|0\rangle \tag{8}$$

where $U(x, \theta)$ is a VQC that depends on the input x and a set of free parameters θ .

In the context of VQC-based Q-learning, considering an environment with $|A|$ possible actions, we need to measure the expectation values of $|A|$ observables, such that the Q-value of a particular action a is given by:

$$Q(s, a) = \langle 0^{\otimes n}|U(s, \theta)^\dagger O_a U(s, \theta)|0^{\otimes n}\rangle \tag{9}$$

where s is the state, n the number of qubits and O_a the observable corresponding to action a .

The observables used throughout this work depend on the ansatz and number of qubits, but always consist of single-qubit σ_z ’s or tensor products of σ_z ’s applied to different qubits, similar to what has been previously used in the literature (Skolik et al. 2022).

Returning to the general scenario of a function $f_\theta(x)$, the cost function $L(\theta)$ is a function of the expectation value itself. In many scenarios, $L(\theta)$ reduces to $f_\theta(x)$. However, in the context of machine learning, classical post-processing is usually added to the expectation value. For instance, consider the typical supervised learning scenario. Given a labeled dataset $D = \{(x_i, y_i)\}^M$, the typical cost function is the Mean Squared Error (MSE), which can be expressed as follows:

$$L(\theta) = \frac{1}{M} \sum_{i=0}^{M-1} (f_\theta(x_i) - y_i)^2 \tag{10}$$

In practice, the value of f_θ for a specific input x is given by running the VQC multiple times and averaging over the




Fig. 2 Subfigure (a) Skolik’s Architecture. When Data Re-Uploading is used, the whole circuit is repeated in each layer. Otherwise, just the part that is not surrounded by dashed lines. Subfigure (b) UQC Architecture. Each processing layer U is given by $U^{UAT}(\vec{x}; \vec{\omega}, \alpha, \varphi) =$

$R_y(2\varphi)R_z(2\vec{\omega} \cdot \vec{x} + 2\alpha)$ and $\vec{\theta}_i = (\vec{\omega}, \alpha, \varphi)$. Although a single-qubit ansatz was shown for simplicity, this ansatz can be generalized to allow multiple qubits

results. Then, typically, gradient-based methods are used to update the parameters of the VQC. However, gradient-free methods also exist and have their advantages and disadvantages, see Cerezo et al. (2021).

When using gradient-based methods, as will be the case of this work, the gradient of the cost function with respect to the parameters is calculated using the Parameter-Shift Rules (Schuld et al. 2019). These rules state that the partial derivative of $f_{\theta}(x)$ w.r.t a single variational parameter (assuming this parameter is the angle of a Pauli-rotation) is given by:

$$\frac{\partial f_{\theta}(x)}{\partial \theta_i} = \frac{1}{2} [f_{\theta}(x; \theta_i + \pi/2) - f_{\theta}(x; \theta_i - \pi/2)] \quad (11)$$

To compute the partial derivative of the cost function w.r.t a single parameter, two circuit executions are needed. Thus, to compute the gradient of the cost function, which has p parameters, $2 \times p$ executions are necessary. However, it is important to note that, in the numerical experiments of this paper, quantum simulators and the adjoint differentiation techniques will be used (Jones and Gacon 2020).

3.4 VQC’s trainability

Let $L(\theta)$ be a hardware-efficient VQC cost function. The gradient of this function is given by:

$$\nabla_{\theta} L(\theta) = \left\{ \frac{\partial L}{\partial \theta_1}, \frac{\partial L}{\partial \theta_2}, \dots, \frac{\partial L}{\partial \theta_k} \right\}, \theta \in \mathbb{R}^k \quad (12)$$

where θ are the cost function’s parameters and k the number of parameters.

The Barren Plateau Phenomenon asserts that

$$Var \left(\frac{\partial L(\theta)}{\partial \theta_k} \right) \in \mathcal{O} \left(\frac{1}{\alpha^n} \right), \alpha > 1 \quad (13)$$

where n is the number of qubits.

Put in words, the variance of the gradient of hardware-efficient VQC-based cost functions decays exponentially

with the number of qubits (McClean et al. 2018). Furthermore, Holmes et al. (2022) demonstrated that the more expressive the ansatz, the lower the variance in the cost gradient. Consequently, the cost landscape is flatter, making these circuits extremely hard to train. Nonetheless, it is important to emphasize that Cerezo et al. (2021) observed that the Barren Plateau Phenomenon highly depends on the locality of the observable employed. To mitigate this issue, several methods have been proposed, such as Barren-Plateau Free ansatzes (Park and Killoran 2024; Schatzki et al. 2024). However, Cerezo et al. (2023) claim that circuits that can be proved to be free from Barren Plateaus actually end up encoding the problem in a polynomial sub-space of the Hilbert Space, one that is classically simulatable.

It is important to note that the BP Phenomenon is a statement over the landscape on average (Drudis et al. 2024), thus it does not exclude the possibility of there being fertile valleys — regions with high gradients near good local minima. Research on warm starts aims to initialize the VQCs in such regions of the landscape (Mele et al. 2022; Zhang et al. 2022; Grant et al. 2019). Care must be taken, however, since regions with large gradients are not necessarily near a good local minima. In fact, it has been shown that hardware-efficient VQCs’ landscapes are swamped with bad local minima [46].

3.5 Data re-uploading

The authors of Schuld et al. (2021) determined that the data encoding strategy used in the quantum circuit influences the expressive power of the quantum model. Specifically, they show that a VQC may be written as a Partial Fourier Series, where the accessible data frequencies are determined by the eigenvalues of the data encoding Hamiltonians. Furthermore, quantum models can access increasingly rich frequency spectra by repeating simple data encoding gates multiple times either in series or in parallel, a technique called *Data Re-uploading* (Pérez-Salinas et al. 2020).

To exemplify the power of data re-uploading, the authors showed that a quantum model employing a single Pauli-

rotation encoding can only learn a sine function of the input. In other words, if the encoding is done just once, the VQC can only approximate functions with a single non-zero frequency in the frequency spectrum. However, repeating Pauli encoding linearly extends the frequency spectrum, such that, if the encoding is repeated n times, then the quantum model can approximate functions with up to n different non-zero frequencies.

3.6 VQC-based deep q-learning

Just like NNs can be used to approximate either a value function or a parameterized policy, so can VQCs. The result is a hybrid classical-quantum algorithm that generally works as follows. The agent observes some state s_t and applies some classical pre-processing ϕ . Then, the result $\phi(s)$ is encoded into a VQC $U_\theta(\phi(s))$ using some data encoding technique. The VQC, with the current parameters θ_t , prepares a quantum state. An observable O_a is measured for each possible action. The expectation values of these observables $\langle O_a \rangle_{s,\theta}$ are post-processed and the result represents either the Q-values $Q_\theta(s, a)$ in value-based methods or the policy $\pi_\theta(a|s)$ in policy-based methods. Then, the agent chooses an action a_t using these predictions and executes it in the environment. The reward r_t and the consecutive state s_{t+1} are observed by the classical optimizer. Using the parameter-shift rule, the gradients of the VQC w.r.t the parameters θ_t are calculated. Finally, the classical optimizer determines the new parameters θ_{t+1} (Meyer et al. 2022).

In VQC-based Deep Q-Learning, in order for the Q-values estimated by the quantum circuit to match the optimal Q-values in the environment, the expectation values of the observables will be multiplied by a classical trainable weight, a technique known as *Trainable Output Scaling* (Skolik et al. 2022).

To the best of our knowledge, Chen et al. (2020) was the first paper to use VQCs as function approximators in RL algorithms. Specifically, the authors used VQCs to approximate the Q-function in a Deep Q-Learning algorithm. The pseudo-code for the algorithm can be seen in Algorithm 2.

4 Numerical results

In this section, we present and analyze experimental numerical results concerning both the performance and trainability of VQC-based DQN agents. The results are summarized in Table 1.

For the Skolik models, the input and output scaling weights were initialized as 1 s while the rotational parameters $\theta \sim U[0, \pi]$. For the UQC models, the weights $\vec{w} \sim \mathcal{N}(0, 0.01)$, the bias \vec{b} were initialized as 0s, the rotational parameters $\varphi \sim U[0, \pi]$ and the output scaling weights as

Algorithm 2 Variational Quantum Deep Q-Learning.

```

1: Algorithm parameters: step size  $\alpha \in [0, 1]$ , small  $\epsilon \in [0, 1]$ 
2: Initialize replay buffer  $\mathcal{D}$  to capacity  $N$ 
3: Initialize action-value function  $Q$  (quantum circuit) with random parameters  $\theta$ 
4: Initialize target action-value function  $\hat{Q}$  (target quantum circuit) with parameters  $\theta^- = \theta$ 
5: for episode= 1,  $M$  do
6:   Initialize  $s_1$  and encode into the quantum state
7:   for  $t = 1, T$  do
8:     With probability  $\epsilon$  select a random action  $a_t$ , otherwise select
        $a_t = \max_a Q^*(s_t, a; \theta)$  from the output of the quantum circuit
9:     Execute action  $a_t$  in emulator and observe  $r_{t+1}, s_{t+1}$ 
10:    Store experience  $(s_t, a_t, r_{t+1}, s_{t+1})$  in replay memory  $\mathcal{D}$ 
11:    Sample random mini-batch of transitions  $(s, a, r, s')$  from  $\mathcal{D}$ 
12:    Compute the TD-targets using target quantum circuit (with old, fixed parameters  $\theta^-$ )
13:    Perform a gradient descent step on
        $\left[ r + \gamma \max_{a'} \hat{Q}(s', a'; \theta_i^-) - Q(s, a; \theta_i) \right]^2$  w.r.t parameters  $\theta$ 
14:    Every  $C$  steps update  $\hat{Q} = Q$ 
15:    Initialize  $s_t$  and encode into the quantum state
16:   end for
17: end for

```

1 s. For the full hyperparameters for each experiment, see Section 6.

4.1 Methodology for analyzing performance and trainability

In any RL problem, the goal is for the agent to find the optimal or a near-optimal policy, which translates to the behavior that accumulates the most rewards. Consequently, performance is typically measured by the return models achieved in benchmark environments.

To analyze performance, we use the following methodology: N agents are randomly initialized according to the model. These agents are then trained in the environment over M episodes. Should an agent solve the environment before completing the M episodes, its training ceases (i.e., no further updates to the model's parameters), but the agent continues interacting with the environment using its learned knowledge until all episodes conclude. The returns from each episode are collected for all agents. Ultimately, we average the returns across all N agents for each episode, compute the standard deviation, and plot the results. The best-performing model will be the one that achieves the highest average return in less episodes.

To analyze trainability, the following methodology is used: when analyzing a model's performance, we also store the gradient of the cost function with respect to the parameters at each training step, which is every step where the parameters are updated. Subsequently, we compute the norm of these gradients for every training step. To mitigate the effects of stochasticity, we calculate the average norm of the gradients

Table 1 Experimental numerical results

Figure	Explanation	Conclusions
Figure 3	Analysis of the effect data re-uploading and input/output scaling have on the performance of QRL agents in CartPole-v0 and Acrobot-v1 environments	Data re-uploading, input scaling, and output scaling all increase the agents' performance
Figure 4	Analysis of the trainability of the agents from Fig. 3	Gradients tend to increase throughout training and Data Re-uploading models exhibit higher gradients than baseline models
Figures 5 and 6	Analysis of the loss and its gradients throughout training for different values of the target network update frequency	It's the instability that arises from low target network update frequency values that leads to the gradients' behavior observed in Fig. 4
Figure 7	Performance of the Single and Multi-Qubit UQC with and without entanglement on the Cartpole-v0 and Acrobot-v1 environments	When no entanglement is used, all the features of the state must be encoded into all qubits, otherwise the agent's lack the necessary information to achieve (near) optimal policies
Figure 8	Variance of the gradient for 1000 uniformly initialized UQC agents with local and global cost functions as the number of qubits increases	The variance of the gradient decays exponentially when a global cost function is used and in the border of a polynomial to exponential regime when a local cost function is used
Figure 9	Performance and norm/variance of the gradients throughout training of the UQC models for an increasing number of qubits	The gradients exhibit a similar behavior to what was previously seen. Moreover, as the number of qubits increases, these gradients remain high
Figures 10 and 11	Performance and Gradients of the UQC and Skolik models trained on a binary classification dataset for an increasing number of qubits	The UQC model consistently achieves higher accuracies than the Skolik models. Moreover, these accuracies do not decrease as the number of qubits increases. Finally, the gradients do not exhibit the same behavior as seen in Fig. 9. Thus, the moving targets of Deep Q-Learning contribute to the large gradients found throughout training

across the N agents and the variance of these norms at each training step. However, since agents cease training once the environment is solved, different agents have a different number of training steps. As a result, we limited our procedure to only encompass the steps up until the first agent successfully solved the environment. To facilitate the visualization of the results, a rolling average of the last 100 training steps is used for both the norm of the gradients and the variance of the norms.

4.2 Environments

OpenAI Gym, an open-source Python library, offers a standardized API that bridges learning algorithms with a standard set of environments, providing a platform to analyze and compare RL algorithms. Two specific environments were chosen: CartPole-v0 and Acrobot-V1. The preference for these is twofold: firstly, they are established benchmark environments that have been frequently adopted in academic research for algorithm testing and comparison. Secondly, these environments strike a balance in complexity; they present sufficient challenges to test the robustness of VQC-based algorithms, yet remain tractable for experimental

purposes. While CartPole has already been the subject of investigation in several studies, Acrobot presents a heightened degree of complexity and, intriguingly, hasn't yet been solved using VQC-based Q-Learning.

CartPole consists of a cart that moves on a frictionless track with a pendulum placed upright. The goal is to balance the pole by applying forces in the left and right direction on the cart. The state space is composed of 4 features that include the cart's position and velocity and the pole's angle and angular velocity. The reward function is very simple: the agent receives a +1 reward per time step and there are a maximum of 200 time steps.

Acrobot consists of two links that are linearly connected to form a chain, with one end of the chain fixed. The joint between the two links is the one we can act upon and the goal is to apply torques (-1,0 or 1) to this joint to swing the free end of the chain above a given height while starting from the initial position of hanging downwards. The state space is more complex than Cartpole's and consists of 6 features. Nonetheless, we employed a trivial technique to reduce it to 4. Every step that does not reach the goal receives a reward of -1 and achieving it results in termination with a reward of 0. Moreover, the maximum number of steps in an episode is 500, after which the episode also terminates.

4.3 Data re-uploading’s effect on performance

To illustrate the impact of data re-uploading and trainable input and output scaling, we analyzed the performance of models composed of all the possible combinations of these characteristics, see Fig. 3.

Figure 3a and b show the considerable effects of data re-uploading and trainable input and output scaling on the performance of VQC-based Deep Q-Learning agents in the CartPole-v0 and Acrobot-v1 environments. Agents that do not use trainable output scaling perform very poorly, no better than random guessing, demonstrating the importance of matching the range of the Q-values estimated by the quantum circuit to the range of the optimal Q-values of the specific environment. While it is possible to achieve a considerable level of performance by multiplying the expectation values by a fixed value instead of using a classical trainable weight,

Skolik et al. (2022) show that such a technique depends on the speed of convergence of the models and does not achieve the same level of performance as models that use trainable output scaling.

Moreover, the impact of trainable input scaling is noticeable, considering that models that use it consistently outperform models that do not. This result emphasizes the importance of allowing for an adaptive frequency matching between the function the VQC outputs and the one it attempts to approximate (Pérez-Salinas et al. 2020).

Finally, most models that use data re-uploading outperform those that do not in both environments, which reaffirms the results obtained by Skolik et al. (2022), demonstrating the importance of such a technique in increasing the expressivity of the VQC and, consequently, allowing for a better approximation of the optimal Q-function. However, this is not the case for the models without trainable input scaling in

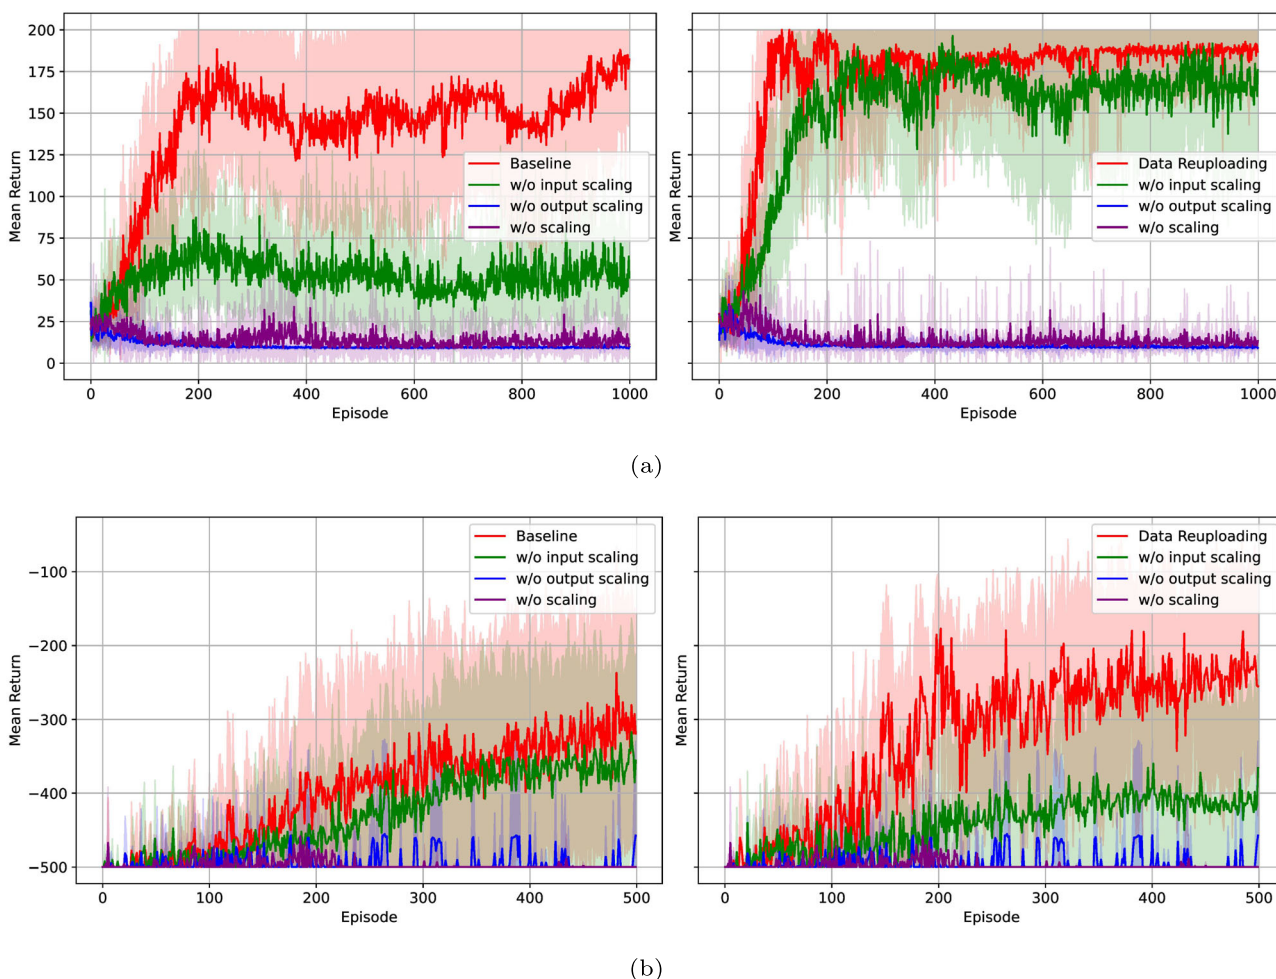


Fig. 3 Performance of Baseline Models (on the left) and Data Re-Uploading models (on the right) in the CartPole-v0 environment (see subfigure (a)) and the Acrobot-v1 environment (see subfigure (b)) with

and without trainable input and/or output scaling. The returns are averaged over 10 agents. The full set of hyperparameters can be seen in Table 3

the Acrobot-v1 environment, which is surprising. This result may possibly be explained by the statistical variance inherent to our methodology for analyzing performance. Nonetheless, if the trainability of data re-uploading models is considerably reduced, then the applicability of data re-uploading might be severely limited. Hence, the next section analyses the trainability of these models.

4.4 Data re-uploading’s effect on trainability

In this section, we turn our attentions to the trainability of the models from Fig. 3a and b. Since the models without trainable output scaling achieved a poor level of performance, they were excluded from this analysis. The results can be seen in Fig. 4a and b.

These results are noteworthy in several ways. First, there is a pattern in the gradient’s norm and variance. From Fig. 4a, one can see that they increase in the training’s early stages,

peak and then decrease on Cartpole-v0. Interestingly, they increase when the agent needs to learn the most (when it is initialized and the environment is new) and start to decrease around the 10,000 training step mark when the agent should have done most of the learning already and started converging to a policy. However, the gradients of the models trained on Acrobot-v1 consistently increase throughout training, see Fig. 4b. This could possibly be explained by the fact that the agents never reached a consistently high return in this environment.

Moreover, despite their greater circuit depth and enhanced expressivity, the data re-uploading models exhibit the highest gradient’s variance among all models. This finding is in stark contrast to our initial expectations derived from Holmes et al. (2022), which suggests a tradeoff between expressivity and trainability. We should nonetheless note that Holmes et al. (2022) claim that more expressive models are harder to train in the context of Barren Plateaus, that is, the gradient’s variance of such models decreases exponentially with

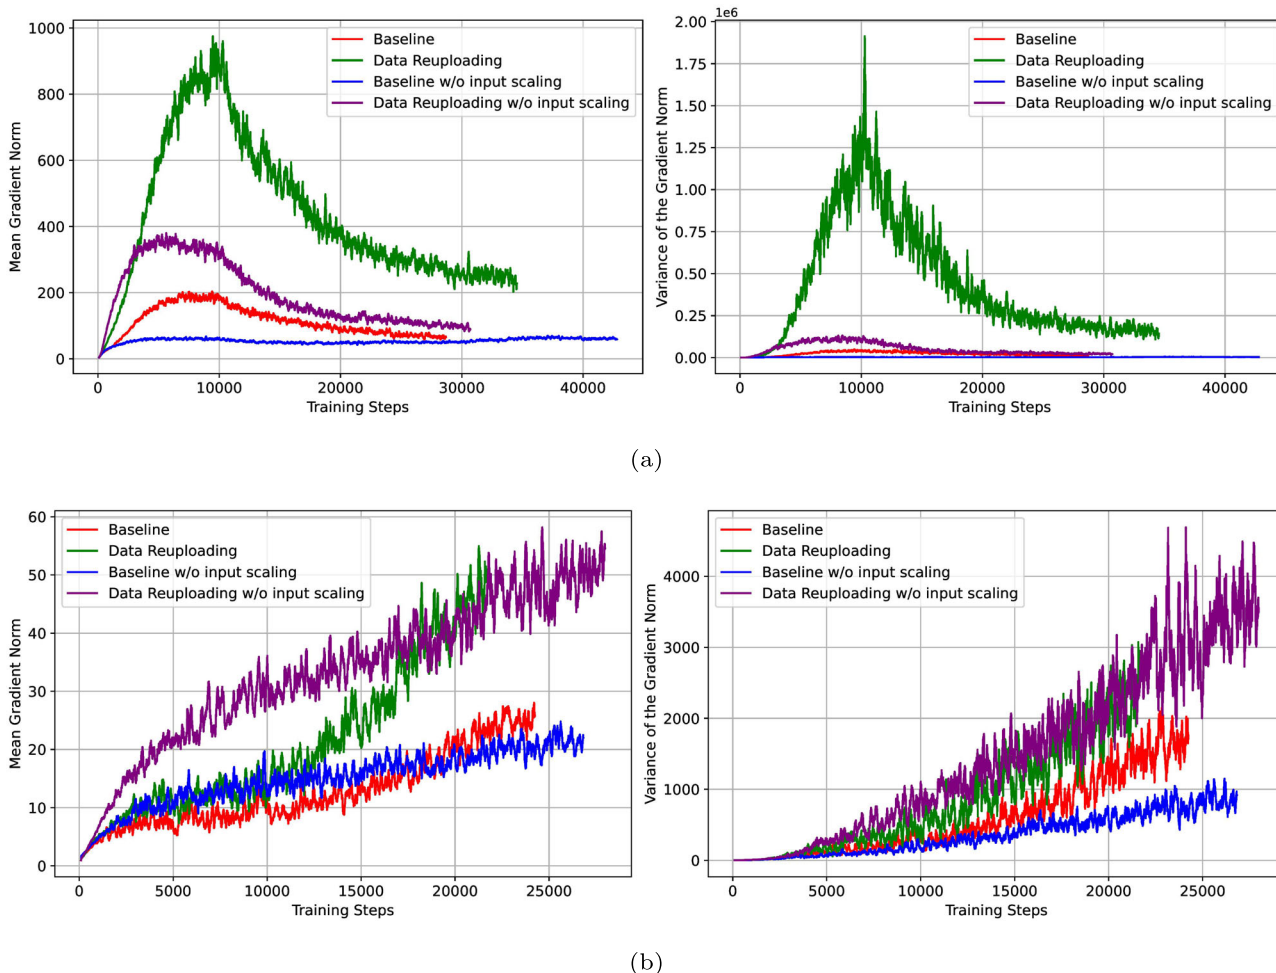


Fig. 4 Trainability of the models from Fig. 3a and b with trainable output scaling in the CartPole-v0 (see subfigure (a)) and Acrobot-v1 (see subfigure (b)) environments. In both subfigures, the left graph represents the gradient’s norm throughout training and the right graph the variance of the norm

system size upon initialization (when parameters are initialized uniformly spanning the whole Hilbert Space). Thus, our results and theirs are not mutually exclusive in the sense that our models may still suffer from the same problem upon uniform initialization covering the whole Hilbert Space, but our initialization in the context of Deep Q-learning leads to large gradients throughout training, a similar concept to warm starts.

4.5 Tradeoff between moving targets and gradient magnitude

In this section, we attempt to comprehend why the norm and the variance of the gradients exhibit such a surprising behavior, both in the sense that they reach substantial maximum values and in the sense that more expressive models achieve higher maximum values. For a start, it is crucial to understand that Deep Q-Learning is fundamentally different from supervised learning, with the primary reason being that Deep Q-Learning targets are *non-stationary*. Since the targets keep changing during training due to the agent's evolving knowledge, predicting Q-values becomes increasingly challenging. This is specially pronounced in the beginning of training, when the agent is focused on exploring the state space. The more states explored, the higher the variance in the return and, consequently, the higher the loss.

Due to the instability that arose from the moving targets, the original paper (Mnih et al. 2015) introduced the concept of a target network. The targets are calculated using this network that has frozen weights, which are updated every C steps to match the weights of the online network. Hence, for C steps at a time, the targets appear stationary. If C is set too high, the targets move slowly and the model takes longer to train. If C is too low, then the targets change frequently and the algorithm becomes unstable.

Figure 5a and b show how data re-uploading models perform for different values of C and the behavior of their loss functions in the CartPole-v0 and Acrobot-v1 environments.

As C increases to substantial values, the models' speed of convergence and performance start to decrease, as expected. However, interestingly, really low values of C (even a value of $C = 1$, which corresponds to a model that does not use a target network), achieve a level of performance equal to or even better than moderate values of C . In fact, after a search of hyperparameters, Skolik et al. (2022) found that $C = 1$ leads to the best performance. Turning our attention to the loss functions, we see that, as C increases, the loss function becomes more stable and the maximum values it achieves start decreasing. Moreover, the dynamics of the moving targets become more evident, as one can see by the loss function for $C = 2500$ in the Cartpole-v0 environment, where the peaks correspond to the step the targets change,

and the valleys in between them to the steps for which they remain stationary.

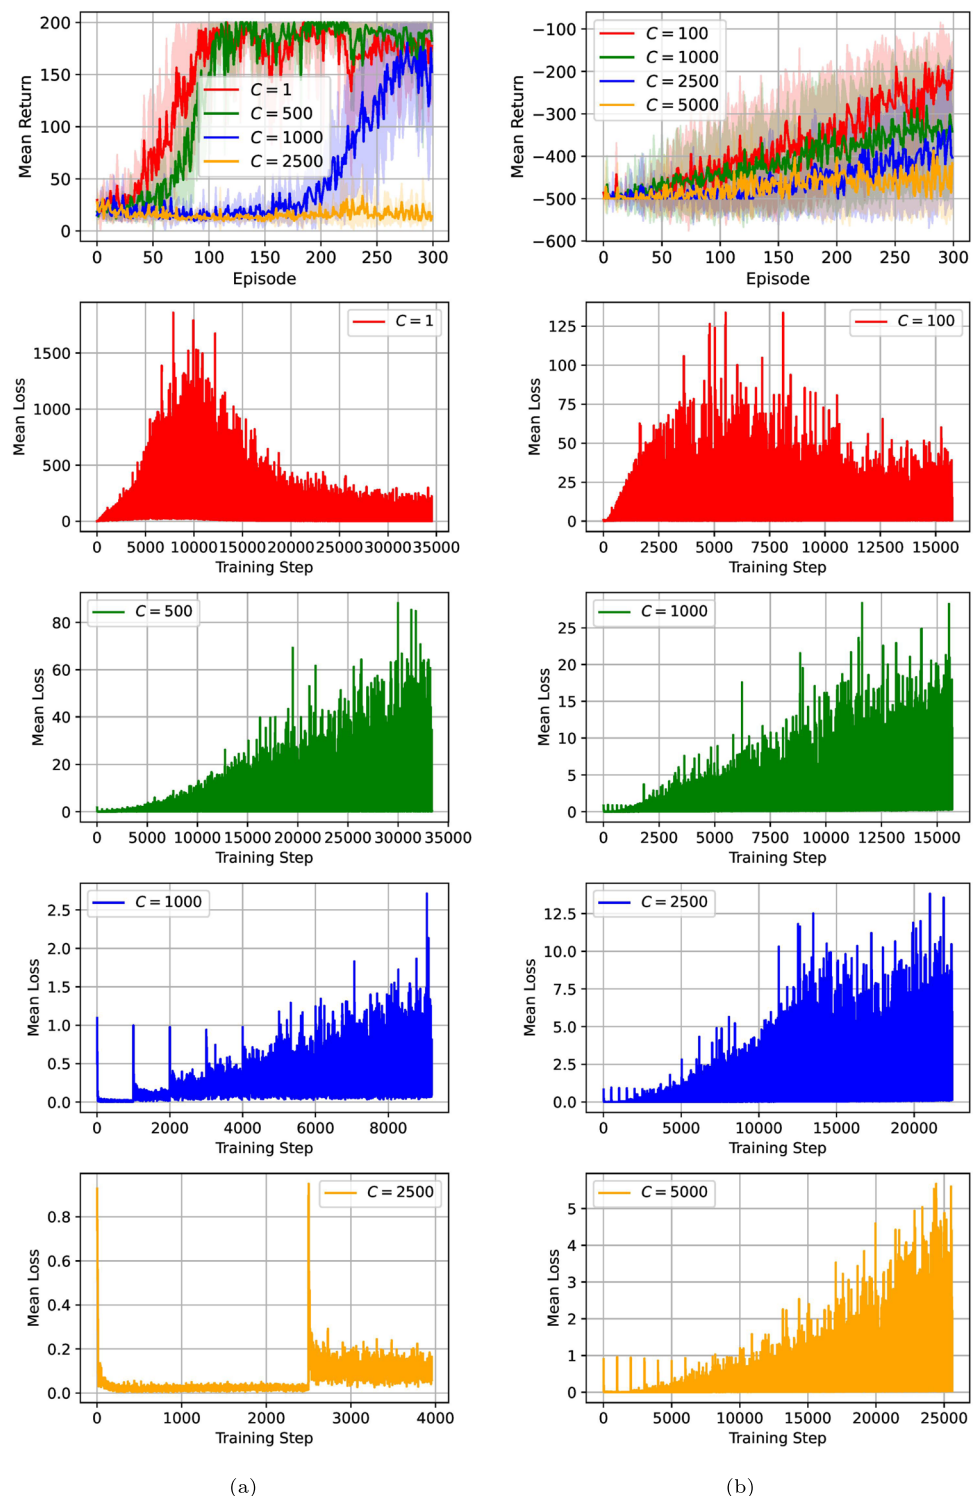
Interestingly, both the gradients' norm and variance exhibit a similar behavior to the loss functions, see Fig. 6. This analysis is of utmost importance for VQC-based Deep Q-Learning. In the previous sections, we observed that well-performing models exhibit substantial gradient magnitudes and variances throughout training. This section has empirically demonstrated the influential role of moving targets and the target network update frequency C on the behavior of both the loss function and, consequently, the gradients. In particular, we saw that increasing C stabilizes the loss function, which contributes to more controlled magnitudes of gradients and their variance. However, it is important to highlight that even with relatively low C values — where the gradients' magnitudes and variances were pronounced — the models showcased impressive performance. They were capable of achieving maximum returns in as few episodes as other more "stable" models with higher values of i .

This raises an interesting question. It is well known that hardware-efficient VQCs suffer from various trainability issues, from optimization landscapes swamped with bad local minima (Anschuetz and Kiani 2022) to the Barren Plateau Phenomenon (McClellan et al. 2018), such that the variance of the gradients decays exponentially with system size. However, it appears that Deep Q-Learning is an inherently unstable algorithm due to the moving targets. Intriguingly, certain hyperparameter configurations seem to allow Deep Q-Learning models to effectively learn good policies while sustaining considerable gradient magnitudes and variances. This hints at a potential advantage: might this inherent instability help counteract the aforementioned trainability issues? We have seen that the gradient's variance of the quantum models increases throughout training. Thus, even if these same models suffer from the Barren Plateau Phenomenon, it may be possible to train them in the context of Deep Q-Learning due to the large gradients seen throughout training. To verify these claims, it becomes imperative to research the gradient's behavior for an increasing number of qubits. The most appropriate architecture for such a study is the UQC, since it allows the encoding of an input vector into an arbitrary number of qubits, leading to a higher flexibility in the choice of system size. Thus, in the next section we introduce this architecture and, then, in Section 4.7, we analyze how the gradients behave as the number of qubits increases.

4.6 The single and multi-qubit UQC

We verified in Section 4.3 that the VQC used in Skolik et al. (2022) is capable of solving the CartPole-v0 environment and achieving considerable performance in the Acrobot-v1 environment. However, it would be interesting to see if the

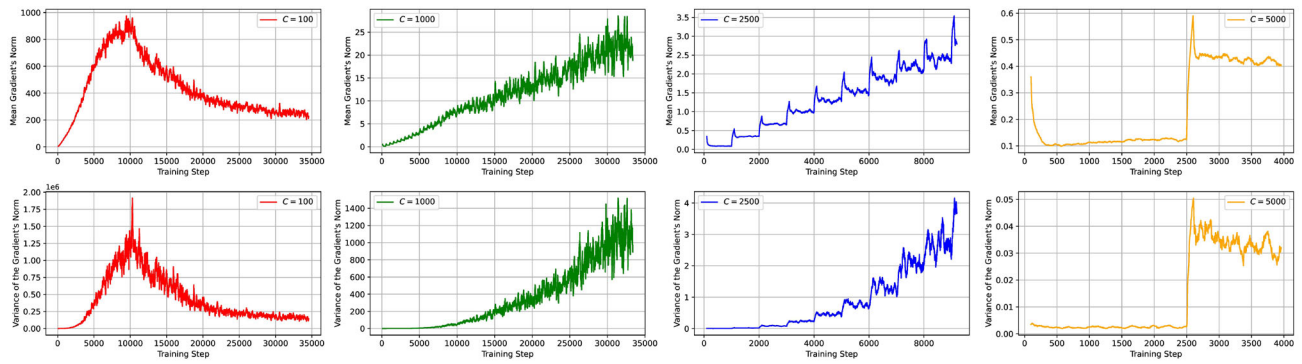
Fig. 5 Performance (first graph) and the respective loss functions for increasing values of C of data re-uploading models in the Cartpole-v0 (see subfigure (a)) and Acrobot-v1 (see subfigure (b)) environments. The full set of hyperparameters can be seen in Table 4



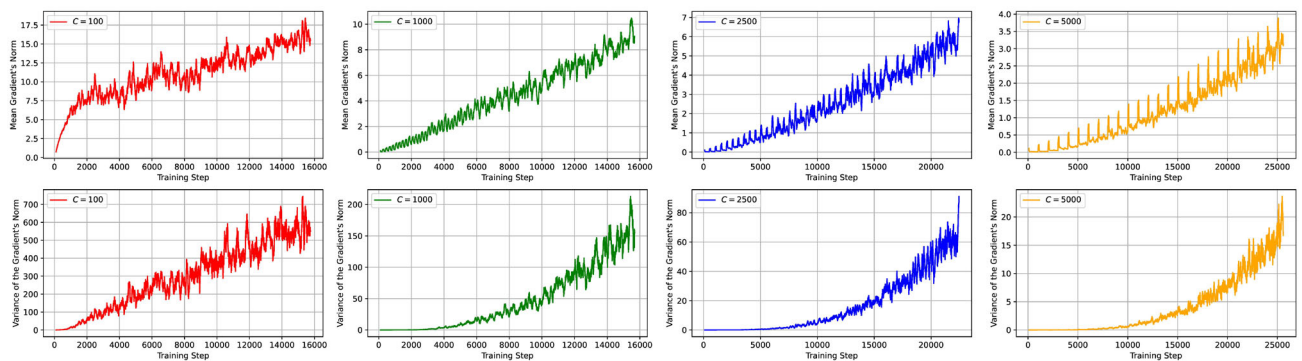
UQC (Pérez-Salinas et al. 2020), a circuit that also makes use of data re-uploading and trainable input scaling by default, is capable of solving these environments. An advantage of this architecture is the fact that it allows the encoding of any input vector using an arbitrary number of qubits. This raises some

interesting questions: Can a single-qubit UQC solve these environments? By increasing the number of qubits, can we improve the performance of the models?

To start, a data encoding technique has to be defined. We experiment with two different types of encoding:




(a)




(b)

Fig. 6 Gradients' norm (above) and variance (below) for data re-uploading models with increasing values of C in the Cartpole-v0 (see subfigure (a)) and Acrobot-v1 (see subfigure (b)) environments

Fig. 7 Performance of UQC models with entanglement (on the left) and without entanglement (on the right) for the CartPole-v0 (see subfigure (a)) and the Acrobot-v1 (see subfigure (b)) environments. The returns are averaged over 10 agents. The full set of hyperparameters can be seen in Table 5



(a)



(b)




Fig. 8 Variance of the Gradient’s Norm across 1000 uniformly sampled parameters spanning the entire Hilbert Space. Two cost functions are tested, a global cost function and a local cost function. For the full set of hyperparameters, see Table 7

- Full Encoding:** Using full encoding, the whole input vector is encoded into all the qubits. Consequently, the number of parameters grows linearly with the number of qubits.

- Partial Encoding:** In this data encoding technique, we divide the number of the input vector features by the number of qubits used and encode a different subvector in each qubit.

The reason for using these two types of encoding is twofold. On the one hand, these two techniques allow us to study the impact of introducing entanglement on the performance of the models in the CartPole-v0 environment. On the other hand, the Full Encoding technique allows us to encode a given input vector in an arbitrary number of qubits that may even be greater than the number of features of the input vector. Thus, we may study how the performance and trainability behave as the number of qubits increases, see Section 4.7. Figure 7a and b show how each of these data encoding methods perform with and without entanglement for 2 and 4 qubits in the Cartpole-v0 and Acrobot-v1 environments, see Fig. 7a and b, respectively.

Models without entanglement, that is, models that can be efficiently simulated using classical computers, are able of achieving considerable performance. Moreover, while models with full encoding are capable of learning (sub-)optimal policies, models that use partial encoding require entanglement, such that the features of the input state become correlated, giving the VQC the full information necessary to better approximate the optimal Q-values of the two envi-

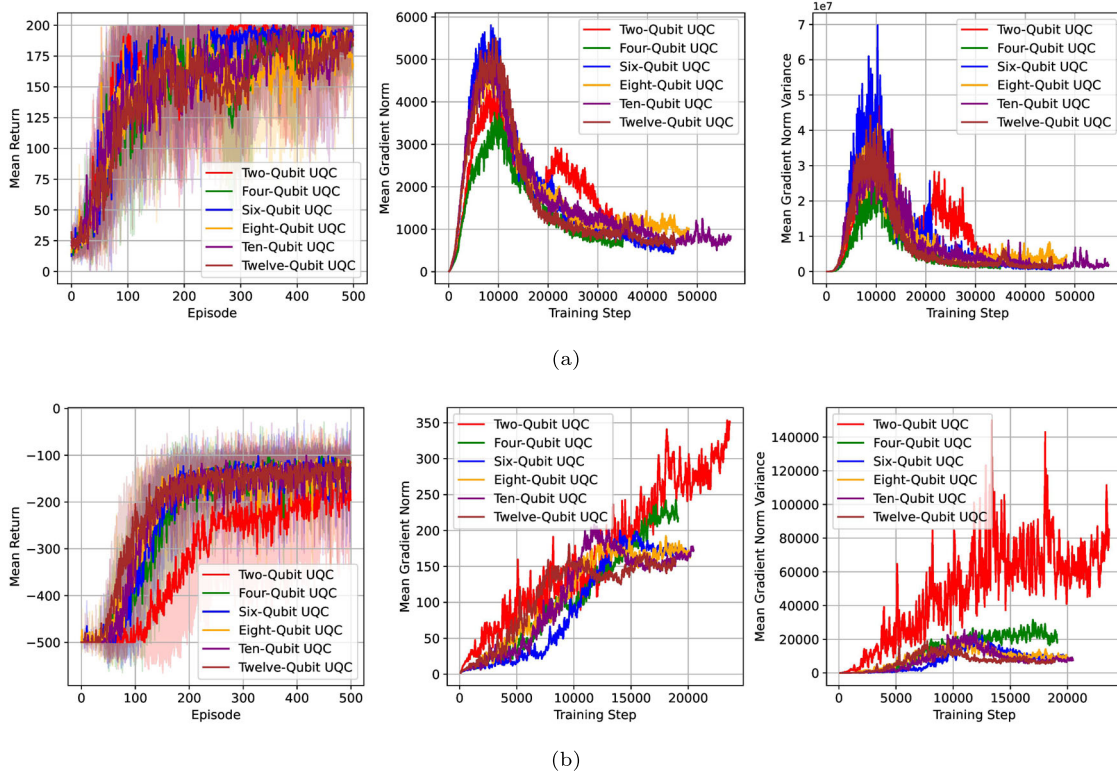


Fig. 9 Performance (first graph) and Trainability Analysis UQC models with increasing numbers of qubits in the Cartpole-v0 (see subfigure (a)) and Acrobot-v1 (see subfigure (b)) environments. The full set of hyperparameters can be seen in Table 6

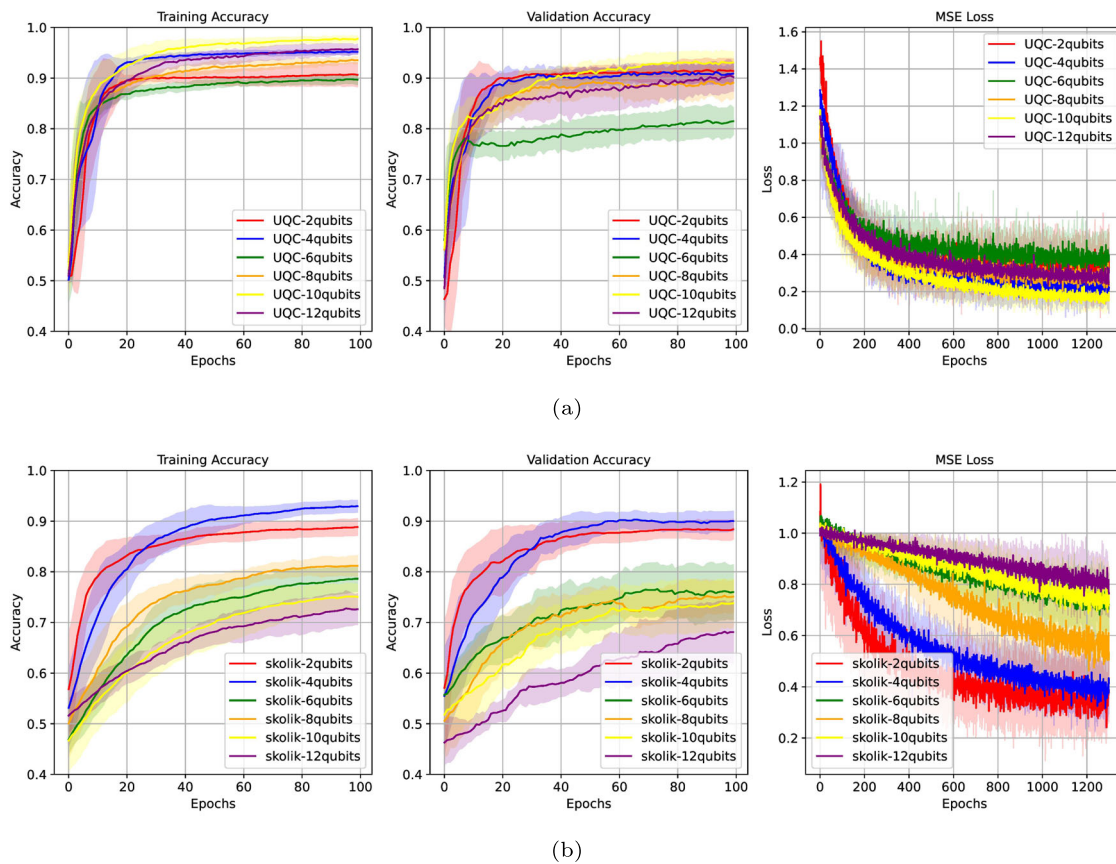


Fig. 10 Training and Validation accuracy as well as MSE cost function throughout training for the Skolik Data Re-Uploading (subfigure (a)) and Full-Encoding Multi-Qubit UQC (see subfigure (b)) trained on a binary classification problem. The full set of hyperparameters can be seen in Table 8

ronments. Finally, even a single-qubit UQC is capable of solving the Cartpole-v0 environment, but the Acrobot-v1 environment needs a VQC with higher width to learn a decent policy.

4.7 Gradient behavior for increasing system sizes

One of the advantages of the multi-qubit UQC architecture is that one may encode any input vector into an arbitrary number of qubits. In particular, it is possible to increase the number of qubits even further than the number of features of the input vector.

We first start by analyzing the gradient's variance of the multi-qubit UQC as the number of qubits increases when a uniform initialization that spans the whole Hilbert Space is used. Thus, we only consider the first training step of CartPole-v0 and plot the variance of the norm of the gradient of 1000 agents for every even number of qubits between 2 and 16, see Fig. 8.

When a global cost function is used, the variance decays exponentially with the number of qubits. However, for a local cost function (the one used throughout this work for the QRL agents), it seems to decay in a regime on the border of polynomial to exponential (as indicated by the R2 error on the fit), a result consistent with Cerezo et al. (2021).

However, we are particularly interested to see how the gradients behave throughout training when using fine-tuned initialization strategies. Thus, we trained the full encoding multi-qubit UQC model with the set of even numbers of qubits from 2 to 12 in the CartPole-v0 and Acrobot-v1 environments, see Fig. 9a and b, respectively.

Both the magnitude of the gradient and its variance behave similarly during training for all models, while also achieving a similar range of values. However, it is important to note that there seems to be a decrease in the variance of the gradients for the Acrobot-v1 environment as the number of qubits increases, but this seems to be most noticeable after 10,000 training steps, after which most of the training should be over. These results are particularly interesting

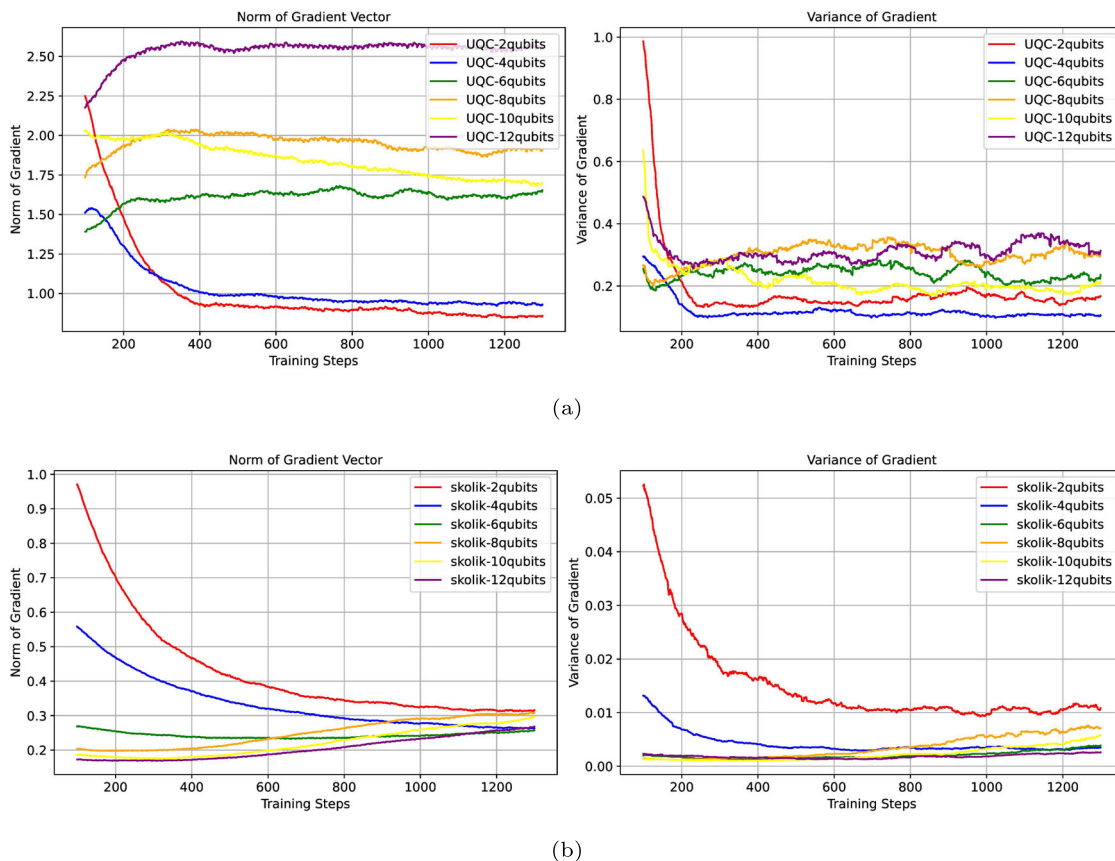


Fig. 11 Norm and Variance of the norm of the gradient vector throughout training for the Skolik Data Re-Uploading (subfigure (a)) and Full-Encoding Multi-Qubit UQC (see subfigure (b)) trained on a binary classification problem. The full set of hyperparameters can be seen in Table 8

because they confirm the suspicions that arose throughout this work. Throughout training, all the models exhibit large gradients even as the number of qubits increases. This seems to be similar to warm starts, since even though these landscapes are characterized by exponentially vanishing gradients, the models are able to navigate the landscape in zones with high gradients, facilitating training. This hints at the possibility of hardware-efficient VQCs being especially suitable to be used as function approximators in Deep Q-learning.

4.8 Gradient behavior in supervised learning

So far, we have claimed that the high gradients found throughout training by the VQC-based Deep Q-Learning models are due to the instability introduced by the moving targets. To verify whether this is true, we trained both the Skolik Data Re-uploading and the Full Encoding Multi-Qubit UQC on the more common task of binary classification for an increas-

ing number of qubits. Since the Skolik Architecture requires a number of qubits equal to the number of features of the datapoints, we used `scipy`'s `make_classification` method to generate a dataset of 500 datapoints for each number of qubits tested. To keep this as close to the original RL scenario as possible, we also used the MSE cost function and the same observables as those used previously. The accuracies achieved on the training and validation datasets as well as the MSE cost function throughout training can be seen in Fig. 10

Interestingly, the Multi-Qubit UQC model consistently outperforms the Skolik model, achieving higher training and validation accuracies and a lower cost value. Moreover, this performance does not decrease as the number of qubits increases, unlike what happens to the Skolik Model, where for 12 qubits the validation accuracy does not even reach 70%. However, more than the performance of such models in terms of accuracy, we want to know how the gradients behave throughout training for a comparison with what we

have seen in the VQC-based Deep Q-Learning agents, see Fig. 11.

The Skolik Models showcase a clear trend in both the norm and the variance of the gradients. As the number of qubits increases, both metrics decrease upon initialization and behave similarly throughout training: they first decrease and then converge. However, for the UQC models, there is not a clear trend in the behavior of the norm of the gradient. Moreover, while the variances across the models with different numbers of qubits behave similarly, there is not a clear decrease in the variance unlike what is seen for the Skolik models.

These results show that the gradients behave differently in the context of supervised learning when compared to Deep Q-learning. This observation confers additional support to the hypothesis that it is the instability caused by the moving targets that allow for high gradients throughout training.

5 Conclusion

The main goal of this work was to study the effects of data re-uploading on the performance and trainability of VQC-based Deep Q-Learning models. It was already empirically shown by Skolik et al. (2022) that using data re-uploading increases the performance of these models in the CartPole environment, possibly due to its increased expressivity, which allows the approximation of more intricate functions (Schuld et al. 2021). Throughout this work, we implemented and tested different data re-uploading models in the CartPole-v0 and Acrobot-v1 environments. Our results match those by Skolik et al. (2022) and Schuld et al. (2021), re-affirming the effect of data re-uploading on the performance of the models.

However, a concern would be that the increase in expressivity and circuit depth that arises from the use of data re-uploading could decrease the trainability of these models. To investigate such a concern, we analyzed the norm of the gradients and the variance of this norm. From this analysis, we verified that the gradients achieve substantial values and actually increase when data re-uploading is used versus when it is not used. Furthermore, we empirically showed that this increase is due to an instability that is inherent to Deep Q-Learning, which is the fact that the targets are non-stationary. Alongside the Mean Squared Error loss function, this leads to a very unstable algorithm with substantial gradients. Moreover, we also verified that increasing the number of qubits of the multi-qubit UQC in the CartPole-v0 and Acrobot-v1

environments does not lead to a decrease in the gradients' magnitude or variance throughout training.

5.1 Future work

For future work, it would be interesting to develop a different methodology to analyze the trainability of VQC-based models. The magnitude and variance of the gradients already give some insights into the trainability of a model, but there are other methodologies that could also be used. For instance, it may be possible to study the eigenvalues of the Hessian Matrix.

Another interesting possibility would be to analyze these models from a Fourier Analysis perspective. Schuld et al. (2021) show that VQCs may be seen as Partial Fourier Series in the data and that using data re-uploading increases the frequencies the VQC “has access to”. Thus, it may be possible to analyze the optimal Q-functions of certain environments and the Fourier Series that the VQCs are approximating. This might reveal some insights into what is happening behind the scenes.

Finally, it would be interesting to build upon this work and analyze the trainability of VQC-based Deep Q-Learning models in different, more complex environments, with different loss functions and sets of hyperparameters. It might even be possible to theoretically derive bounds for the gradients of these models, which could reinforce the empirical analysis.

Experiments replication

All data necessary for replicating the results of this paper is available [here](#).

Appendix: Hyperparameters

Some of the hyperparameters for the VQC-based Deep Q-Learning models are explained in the following table:

The set of hyperparameters used for the VQCs in Figs. 3a and b and 4a and b, can be seen in Table 3.

The set of hyperparameters used for the VQCs in Fig. 5a and b can be seen in Table 4.

The set of hyperparameters used for the VQCs in Fig. 7a and b can be seen in Table 5.

The set of hyperparameters used for the VQCs in Fig. 6a and b can be seen in Table 6.

The set of hyperparameters used for the VQCs in Figs. 10 and 11 can be seen in Table 6.

Table 2 An explanation of VQC-Based Deep Q-Learning's hyperparameters

Hyperparameter	Explanation
qubits (n)	The quantum circuit's number of qubits
layers	The quantum circuit's number of layers
γ	The return's discount factor
trainable input scaling	whether trainable input scaling is used or not
trainable output scaling	whether trainable output scaling is used or not
learning rate of parameters θ	the learning rate of parameters θ
learning rate of input scaling parameters	the learning rate of input scaling parameters
Learning rate of output scaling parameters	the learning rate of output scaling parameters
batch size	the batch size
decaying schedule of ϵ -greedy policy	defines the decaying schedule of the ϵ -greedy policy (e.g., exponential)
ϵ_{init}	the initial value of ϵ
ϵ_{dec}	the decay rate of <i>epsilon</i> per episode
ϵ_{min}	the minimum value of ϵ
update model	the update model's frequency
update target model	the target model's update frequency
data re-uploading	whether data re-uploading is used or not
input scaling initialization	how the input scaling parameters were initialized (Skolik Model)
output scaling initialization	how the output scaling parameters were initialized (Skolik Model)
rotational parameters initialization	how the rotational parameters were initialized (θ in Skolik, φ in UQC)
\vec{w} initialization	how the \vec{w} parameters were initialized (UQC)
\vec{b} initialization	how the \vec{b} parameters were initialized (UQC)
observables	observables measured

Table 3 Models' hyperparameters of Section 4.3

	CartPole-v0	Acrobot-v1
qubits (n)	4	4
layers	5	5
γ	0.99	0.99
trainable input scaling	yes, no	yes, no
trainable output scaling	yes, no	yes, no
learning rate of parameters θ	0.001	0.001
learning rate of input scaling parameters	0.001	0.001
learning rate of output scaling parameters	0.1	0.1
batch size	16	32
decaying schedule of ϵ -greedy policy	Exponential	Exponential
ϵ_{init}	1	1
ϵ_{dec}	0.99	0.99
ϵ_{min}	0.01	0.01
update model	1	5
update target model	1	250
size of replay buffer	10,000	50,000
data re-uploading	yes, no	yes, no

Table 3 continued

	CartPole-v0	Acrobot-v1
input scaling initialization	Initialized as 1 s	Initialized as 1 s
output scaling initialization	Initialized as 1 s	Initialized as 1 s
rotational parameters initialization	Uniformly sampled between 0 and π	Uniformly sampled between 0 and π
\vec{w} initialization	—————	—————
\vec{b} initialization	—————	—————
observables	(Z_0Z_1, Z_2Z_3)	(Z_0, Z_1Z_2, Z_3)

Table 4 Models' hyperparameters of Section 4.5

	CartPole-v0	Acrobot-v1
qubits (n)	4	4
layers	5	5
γ	0.99	0.99
trainable input scaling	yes	yes
trainable output scaling	yes	yes
learning rate of parameters θ	0.001	0.001
learning rate of input scaling parameters	0.001	0.001
learning rate of output scaling parameters	0.1	0.1
batch size	16	32
decaying schedule of ϵ -greedy policy	Exponential	Exponential
ϵ_{init}	1	1
ϵ_{dec}	0.99	0.99
ϵ_{min}	0.01	0.01
update model	1	5
update target model	1,500,1000,2500	100,1000,2500,5000
size of replay buffer	10,000	50,000
data re-uploading	yes	yes
input scaling initialization	Initialized as 1 s	Initialized as 1 s
output scaling initialization	Initialized as 1 s	Initialized as 1 s
rotational parameters initialization	Uniformly sampled between 0 and π	Uniformly sampled between 0 and π
\vec{w} initialization	—————	—————
\vec{b} initialization	—————	—————
observables	(Z_0Z_1, Z_2Z_3)	(Z_0, Z_1Z_2, Z_3)

Table 5 Models' hyperparameters of Section 4.6

	CartPole-v0	Acrobot-v1
qubits (n)	1,2,4	1,2,4
layers	5	5
γ	0.99	0.99
trainable input scaling	yes	yes
trainable output scaling	yes	yes
learning rate of parameters θ	0.001	0.001
learning rate of input scaling parameters	0.001	0.001

Table 5 continued

	CartPole-v0	Acrobot-v1
learning rate of output scaling parameters	0.1	0.1
batch size	16	32
decaying schedule of ϵ -greedy policy	Exponential	Exponential
ϵ_{init}	1	1
ϵ_{dec}	0.99	0.99
ϵ_{min}	0.01	0.01
update model	1	5
update target model	1,500,1000,2500	100,1000,2500,5000
size of replay buffer	10,000	50,000
data re-uploading	yes	yes
input scaling initialization	—————	—————
output scaling initialization	Initialized as 1 s	Initialized as 1 s
rotational parameters initialization	Uniformly sampled between 0 and π	Uniformly sampled between 0 and π
\vec{w} initialization	Gaussian Distribution (mean=0, std=0.01)	Gaussian Distribution (mean=0, std=0.01)
\vec{b} initialization	Initialized as 0 s	Initialized as 0 s
observables	(Z_0, Z_1, Z_2, Z_3)	(Z_0, Z_1, Z_2, Z_3)

Table 6 Models' hyperparameters of Section 4.7

	CartPole-v0	Acrobot-v1
qubits (n)	2,4,6,8,10,12	2,4,6,8,10,12
layers	5	5
γ	0.99	0.99
trainable input scaling	yes	yes
trainable output scaling	yes	yes
learning rate of parameters θ	0.001	0.001
learning rate of input scaling parameters	0.001	0.001
learning rate of output scaling parameters	0.1	0.1
batch size	16	32
decaying schedule of ϵ -greedy policy	Exponential	Exponential
ϵ_{init}	1	1
ϵ_{dec}	0.99	0.99
ϵ_{min}	0.01	0.01
update model	1	5
update target model	1,500,1000,2500	100,1000,2500,5000
size of replay buffer	10,000	50,000
data re-uploading	yes	yes
input scaling initialization	—————	—————
output scaling initialization	Initialized as 1 s	Initialized as 1 s
rotational parameters initialization	Uniformly sampled between 0 and π	Uniformly sampled between 0 and π

Table 6 continued

	CartPole-v0	Acrobot-v1
\vec{w} initialization	Gaussian Distribution (mean=0, std=0.01)	Gaussian Distribution (mean=0, std=0.01)
\vec{b} initialization	Initialized as 0s	Initialized as 0s
observables	$(Z_0 \dots Z_{n/2-1}, Z_{n/2} \dots Z_n)$	$(Z_0, Z_1 \dots Z_{n-1}, Z_n)$

Table 7 Models' hyperparameters of Section 4.7 — Fig. 8

	UQC Local	UQC Global
qubits (n)	2,4,6,8,10,12,14,16	2,4,6,8,10,12,14,16
layers	5	5
γ	0.99	0.99
trainable input scaling	yes, no	yes, no
trainable output scaling	yes, no	yes, no
learning rate of parameters θ	0.001	0.001
learning rate of input scaling parameters	0.001	0.001
learning rate of output scaling parameters	0.1	0.1
batch size	16	16
decaying schedule of ϵ -greedy policy	Exponential	Exponential
ϵ_{init}	1	1
ϵ_{dec}	0.99	0.99
ϵ_{min}	0.01	0.01
update model	1	1
update target model	1	1
size of replay buffer	10,000	10,000
data re-uploading	yes	yes
input scaling initialization	—————	—————
output scaling initialization	—————	—————
rotational parameters initialization	Uniformly sampled between 0 and 2π	Uniformly sampled between 0 and 2π
\vec{w} initialization	Uniformly between 0 and 2π	Uniformly between 0 and 2π
\vec{b} initialization	Uniformly between 0 and 2π	Uniformly between 0 and 2π
observables	$(Z_0 \dots Z_{n/2-1}, Z_{n/2} \dots Z_n)$	$(Z_0 \dots Z_n, X_0 \dots Z_n)$

Table 8 Models' hyperparameters of Section 4.8

	UQC	Skolik
qubits (n)	2,4,6,8,10,12	2,4,6,8,10,12
layers	5	5
γ	—————	—————
trainable input scaling	yes	yes
trainable output scaling	yes	yes
learning rate of parameters θ	0.001	0.001
learning rate of input scaling parameters	0.001	0.001
learning rate of output scaling parameters	0.1	0.1
batch size	32	32

Table 8 continued

	UQC	Skolik
decaying schedule of ϵ -greedy policy	—————	—————
ϵ_{init}	—————	—————
ϵ_{dec}	—————	—————
ϵ_{min}	—————	—————
update model	—————	—————
update target model	—————	—————
size of replay buffer	—————	—————
data re-uploading	yes	yes
input scaling initialization	—————	Initialized as 1 s
output scaling initialization	Initialized as 1 s	Initialized as 1 s
rotational parameters initialization	Uniformly sampled between 0 and π	Uniform Sampling between 0 and π
\vec{w} initialization	Gaussian Distribution (mean=0, std=0.01)	—————
\vec{b} initialization	Initialized as 0 s	—————
observables	$(Z_0 \dots Z_{n/2-1}, Z_{n/2} \dots Z_n)$	$(Z_0 \dots Z_{n/2-1}, Z_{n/2} \dots Z_n)$

Author Contributions R.C. implemented the code and wrote the main manuscript text, under the supervision of A.S. and L.P.S. All authors reviewed the manuscript.

Funding This work is financed by National Funds through the Portuguese funding agency, FCT - Fundação para a Ciência e a Tecnologia, within project LA/P/0063/2020 (<https://doi.org/10.54499/LA/P/0063/2020>). RC also thanks the support of the Foundation for Science and Technology (FCT, Portugal) under grant 10053/BII-E_B4/2023. AS also thanks the support of the Foundation for Science and Technology (FCT, Portugal) within grant UI/BD/152698/2022 and project IBEX, with reference PTDC/CCI-COM/4280/2021.

Data Availability All data necessary to replicate the results of this manuscript can be found in https://github.com/RodrigoCoelho7/VQC_Qlearning

Declarations

Competing interests The authors declare no competing interests.

Conflict of interest The authors declare no competing interests.

References

- Andre D, Russell SJ (2002) State abstraction for programmable reinforcement learning agents. In: Aaai/iaai, pp 119–125
- Anschuetz ER, Kiani BT (2022) Quantum variational algorithms are swamped with traps. *Nature Commun* 13(1):7760
- Bilkis M, Cerezo M, Verdon G, Coles PJ, Cincio L (2021) A semi-agnostic ansatz with variable structure for quantum machine learning. arXiv preprint [arXiv:2103.06712](https://arxiv.org/abs/2103.06712)
- Busoniu L, Babuska R, De Schutter B, Ernst D (2017) Reinforcement Learning and Dynamic Programming Using Function Approximators. CRC press, ???
- Cerezo M, Arrasmith A, Babbush R, Benjamin SC, Endo S, Fujii K, McClean JR, Mitarai K, Yuan X, Cincio L et al (2021) Variational quantum algorithms. *Nature Reviews. Physics* 3(9):625–644
- Cerezo M, Larocca M, García-Martín D, Diaz N, Braccia P, Fontana E, Rudolph MS, Bermejo P, Ijaz A, Thanasilp S et al (2023) Does provable absence of barren plateaus imply classical simulability? or, why we need to rethink variational quantum computing. arXiv preprint [arXiv:2312.09121](https://arxiv.org/abs/2312.09121)
- Cerezo M, Sone A, Volkoff T, Cincio L, Coles PJ (2021) Cost function dependent barren plateaus in shallow parametrized quantum circuits. *Nature commun* 12(1):1791
- Chen SY-C, Yang C-HH, Qi J, Chen P-Y, Ma X, Goan H-S (2020) Variational quantum circuits for deep reinforcement learning. *IEEE Access*. 8:141007–141024
- Coyle B, Mills D, Danos V, Kashefi E (2020) The born supremacy: quantum advantage and training of an ising born machine. *npj Quantum Information*. 6(1):60
- Drudis M, Thanasilp S, Holmes Z et al (2024) Variational quantum simulation: a case study for understanding warm starts. arXiv preprint [arXiv:2404.10044](https://arxiv.org/abs/2404.10044)
- Farhi E, Neven H (2018) Classification with quantum neural networks on near term processors. arXiv preprint [arXiv:1802.06002](https://arxiv.org/abs/1802.06002)
- Grant E, Wossnig L, Ostaszewski M, Benedetti M (2019) An initialization strategy for addressing barren plateaus in parametrized quantum circuits. *Quantum* 3:214
- Hambly B, Xu R, Yang H (2023) Recent advances in reinforcement learning in finance. *Math Financ* 33(3):437–503
- Holmes Z, Sharma K, Cerezo M, Coles PJ (2022) Connecting ansatz expressibility to gradient magnitudes and barren plateaus. *PRX Quantum*. 3(1):010313
- Jerbi S, Gyurik C, Marshall S, Briegel HJ, Dunjko V (2021) Variational quantum policies for reinforcement learning. arXiv preprint [arXiv:2103.05577](https://arxiv.org/abs/2103.05577)
- Jones T, Gacon J (2020) Efficient calculation of gradients in classical simulations of variational quantum algorithms. arXiv preprint [arXiv:2009.02823](https://arxiv.org/abs/2009.02823)
- Kaelbling LP, Littman ML, Moore AW (1996) Reinforcement learning: A survey. *J Artif Intell Res* 4:237–285

- Kober J, Bagnell JA, Peters J (2013) Reinforcement learning in robotics: A survey. *Int J Robot Res* 32(11):1238–1274
- LeCun Y, Bengio Y, Hinton G (2015) Deep learning. *nature*. 521(7553):436–444
- Lockwood O, Si M (2020) Reinforcement learning with quantum variational circuit. In: *Proceedings of the AAAI Conference on Artificial Intelligence and Interactive Digital Entertainment*, vol 16. pp 245–251
- McClellan JR, Boixo S, Smelyanskiy VN, Babbush R, Neven H (2018) Barren plateaus in quantum neural network training landscapes. *Nat Commun* 9(1):1–6
- Mele AA, Mbeng GB, Santoro GE, Collura M, Torta P (2022) Avoiding barren plateaus via transferability of smooth solutions in a hamiltonian variational ansatz. *Phys Rev A* 106(6):060401
- Meyer N, Ufrecht C, Periyasamy M, Scherer DD, Plinge A, Mutschler C (2022) A survey on quantum reinforcement learning. arXiv preprint [arXiv:2211.03464](https://arxiv.org/abs/2211.03464)
- Mitarai K, Negoro M, Kitagawa M, Fujii K (2018) Quantum circuit learning. *Phys Rev A* 98(3):032309
- Mnih V, Kavukcuoglu K, Silver D, Rusu AA, Veness J, Bellemare MG, Graves A, Riedmiller M, Fidjeland AK, Ostrovski G, et al (2015) Human-level control through deep reinforcement learning. *nature*. **518**(7540):529–533
- Ostaszewski M, Grant E, Benedetti M (2021) Structure optimization for parameterized quantum circuits. *Quantum* 5:391
- Park C-Y, Killoran N (2024) Hamiltonian variational ansatz without barren plateaus. *Quantum* 8:1239
- Pérez-Salinas A, Cervera-Lierta A, Gil-Fuster E, Latorre JI (2020) Data re-uploading for a universal quantum classifier. *Quantum* 4:226
- Preskill J (2018) Quantum computing in the nisq era and beyond. *Quantum* 2:79
- Sallab AE, Abdou M, Perot E, Yogamani S (2017) Deep reinforcement learning framework for autonomous driving. arXiv preprint [arXiv:1704.02532](https://arxiv.org/abs/1704.02532)
- Schatzki L, Larocca M, Nguyen QT, Sauvage F, Cerezo M (2024) Theoretical guarantees for permutation-equivariant quantum neural networks. *npj Quantum Information*. **10**(1):12
- Schuld M, Killoran N (2019) Quantum machine learning in feature hilbert spaces. *Phys Rev Lett* 122(4):040504
- Schuld M, Bergholm V, Gogolin C, Izaac J, Killoran N (2019) Evaluating analytic gradients on quantum hardware. *Phys Rev A* 99(3):032331
- Schuld M, Sweke R, Meyer JJ (2021) Effect of data encoding on the expressive power of variational quantum-machine-learning models. *Phys Rev A* 103(3):032430
- Schuld M, Bocharov A, Svore KM, Wiebe N (2020) Circuit-centric quantum classifiers. *Phys Rev A* 101(3):032308
- Sequeira A, Santos LP, Barbosa LS (2022) Variational quantum policy gradients with an application to quantum control. arXiv preprint [arXiv:2203.10591](https://arxiv.org/abs/2203.10591)
- Shor PW (1999) Polynomial-time algorithms for prime factorization and discrete logarithms on a quantum computer. *SIAM Rev* 41(2):303–332
- Silver D, Singh S, Precup D, Sutton RS (2021) Reward is enough. *Artif Intell* 299:103535
- Silver D, Huang A, Maddison CJ, Guez A, Sifre L, Van Den Driessche G, Schrittwieser J, Antonoglou I, Panneershelvam V, Lanctot M et al (2016) Mastering the game of go with deep neural networks and tree search. *nature*. **529**(7587):484–489
- Silver D, Hubert T, Schrittwieser J, Antonoglou I, Lai M, Guez A, Lanctot M, Sifre L, Kumaran D, Graepel T et al (2017) Mastering chess and shogi by self-play with a general reinforcement learning algorithm. arXiv preprint [arXiv:1712.01815](https://arxiv.org/abs/1712.01815)
- Skolik A, Jerbi S, Dunjko V (2022) Quantum agents in the gym: a variational quantum algorithm for deep q-learning. *Quantum* 6:720
- Skolik A, Cattelan M, Yarkoni S, Bäck T, Dunjko V (2023) Equivariant quantum circuits for learning on weighted graphs. *npj Quantum Inf* 9(1):47
- Sorzano COS, Vargas J, Montano AP (2014) A survey of dimensionality reduction techniques. arXiv preprint [arXiv:1403.2877](https://arxiv.org/abs/1403.2877)
- Sutton RS, Barto AG (2018) *Reinforcement Learning: An Introduction*. MIT press, ???
- Watkins CJ, Dayan P (1992) Q-learning. *Mach Learn* 8(3):279–292
- Zhang K, Liu L, Hsieh M-H, Tao D (2022) Escaping from the barren plateau via gaussian initializations in deep variational quantum circuits. *Adv Neural Inf Process Syst* 35:18612–18627
- Zoufal C, Lucchi A, Woerner S (2021) Variational quantum boltzmann machines. *Quantum Mach Intell* 3:1–15

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

Springer Nature or its licensor (e.g. a society or other partner) holds exclusive rights to this article under a publishing agreement with the author(s) or other rightsholder(s); author self-archiving of the accepted manuscript version of this article is solely governed by the terms of such publishing agreement and applicable law.