



Robustness of quantum reinforcement learning under hardware errors

Andrea Skolik^{1,2*}, Stefano Mangini^{3,4}, Thomas Bäck¹, Chiara Macchiavello^{3,4,5} and Vedran Dunjko¹

*Correspondence:

a.skolik@liacs.leidenuniv.nl

¹Leiden University, Niels Bohrweg 1,
2333 CA Leiden, The Netherlands

²Volkswagen AG, Ungererstraße 69,
80805 Munich, Germany

Full list of author information is
available at the end of the article

Abstract

Variational quantum machine learning algorithms have become the focus of recent research on how to utilize near-term quantum devices for machine learning tasks. They are considered suitable for this as the circuits that are run can be tailored to the device, and a big part of the computation is delegated to the classical optimizer. It has also been hypothesized that they may be more robust to hardware noise than conventional algorithms due to their hybrid nature. However, the effect of training quantum machine learning models under the influence of hardware-induced noise has not yet been extensively studied. In this work, we address this question for a specific type of learning, namely variational reinforcement learning, by studying its performance in the presence of various noise sources: shot noise, coherent and incoherent errors. We analytically and empirically investigate how the presence of noise during training and evaluation of variational quantum reinforcement learning algorithms affect the performance of the agents and robustness of the learned policies. Furthermore, we provide a method to reduce the number of measurements required to train Q-learning agents, using the inherent structure of the algorithm.

Keywords: Variational quantum algorithms; Quantum machine learning; Quantum hardware noise

1 Introduction

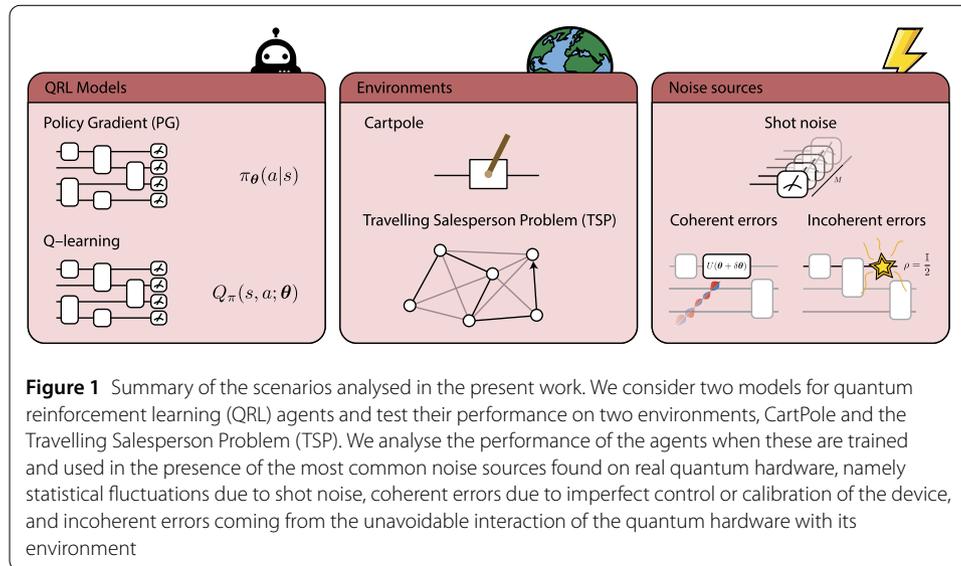
Quantum machine learning (QML) is advertised as one of the most promising candidates for a near-term advantage in quantum computing [1]. The variational quantum algorithms (VQAs) that are used for this are trained in a hybrid fashion, where a classical optimizer is used to tune the parameters of a quantum circuit [2, 3]. It is hypothesized that the hybrid training scheme along with the freedom of adjusting the parameters appropriately, makes these algorithms inherently robust to quantum hardware noise to some extent [2, 4]. This hypothesis is also inspired by classical neural networks, which are robust under certain types of noise. In the classical setting, one can broadly distinguish between two types of noise: benign noise that does not severely impact the training procedure or can even improve generalization [5–8], and adversarial noise which is deliberately constructed to study where neural networks fail [9–12]. Furthermore, we can distinguish between noise that is present during training, and noise that is present when using the trained model. Adversarial noise is usually of the latter case, where a trained neural network can produce

© The Author(s) 2023, corrected publication 2023. **Open Access** This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if changes were made. The images or other third party material in this article are included in the article's Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit <http://creativecommons.org/licenses/by/4.0/>.

completely wrong outputs due to small perturbations of the input data [13]. The benign type of noise mentioned above, on the other hand, is usually present at training time in form of perturbations of the input data, activation functions, weights or structure of the neural network, and has even been established as a method to combat overfitting in the classical literature [5–8, 14]. These results inspired the hypothesis that variational quantum algorithms possess a similar robustness to certain types of noise and may even benefit from its presence when trained on a quantum device. However, thorough investigations that confirm such robustness of VQAs against hardware-related noise, or even a beneficial effect from it, are still lacking. In terms of negative results for the trainability of VQAs under noise, it has been shown that optimization landscapes of noisy quantum circuits become increasingly flat at a rate that scales exponentially with the number of qubits under local Pauli noise when the circuit depth grows linearly with the number of qubits [15]. In the case of the variational quantum eigensolver, where the goal is to find the ground state of a given Hamiltonian, the presence of noise has been shown to lead to increasing deviation from the ideal energy [16]. Similar effects have been studied in the context of the quantum approximate optimization algorithm (QAOA) [17], where the goal is to find the ground state of a Hamiltonian that represents the solution to a combinatorial optimization problem [18, 19].

When it comes to QML, in-depth studies on the effect of noise on the trainability and performance of VQAs are scarce. Apart from the work mentioned above on noise-induced barren plateaus [15], the authors of [20] provided first insights into how the data encoding method used in a quantum classifier influences its resilience to varying types of noise. As for the potential benefit of noise, the authors of [21] show that the stochasticity induced by measurements in a QML model can help the optimizer to escape saddle points. The above results show that, on the one hand, too much noise will make the model untrainable, while on the other hand, modest amounts of noise can even improve trainability [21]. However, it remains unclear how large the gap is between tolerable and harmful amounts of noise [4], and it is not expected that this can be answered in a general way for all different types of learning algorithms and noise sources.

In this work, we shed light on this question from the angle of variational quantum reinforcement learning (QRL). Classical reinforcement learning (RL) models have been shown to be sensitive to noise, either during training [22] or in the form of adversarial samples [23, 24]. Additionally, it is known that a bottleneck of RL algorithms is their sample inefficiency, i.e., many interactions with an environment are needed for training [25]. Still, RL resembles human-type learning most closely among the main branches of modern ML, and therefore motivates further studies in this area. Among these studies, RL with VQAs has been proposed and extensively investigated in the noise-free setting over the past few years [26–34]. These results provide promising perspectives, as quantum models have empirically been shown to perform similarly to neural networks on small classical benchmark tasks [29, 32], while at the same time an exponential separation between classical and quantum learners can be proven for specific contrived environments based on classically hard tasks [28, 29]. These results motivate further studies on how large the above-mentioned gap between tolerable and too much noise is in the case of variational RL algorithms, and how close the algorithm performance can get to the noise-free setting for various types of noise that can be present on near-term devices.



We investigate this for two types of variational RL algorithms, Q-learning and the policy gradient method, by performing extensive numerical experiments for both types of algorithms with two different environments, CartPole and the Travelling Salesperson Problem, and under the effect of a wide class of noise sources, namely shot noise, coherent and incoherent errors. In Fig. 1 we summarise the approach of the present work showing the QRL models, environments and noise sources considered in the analysis. We start by considering the trade-off between the number of measurement shots taken for each circuit evaluation and the performance of variational agents. As the number of shots required by a QML algorithm can be a bottleneck on near-term devices and RL is known to require many interactions with the environment to learn, we propose a method for Q-learning to reduce the number of overall measurements by taking advantage of the structure of the underlying RL algorithm. Second, we model coherent errors with a random Gaussian perturbation of the variational parameters, and analytically study the effect of these perturbations on the output of parameterised quantum circuits, similarly to [35]. We provide an upper bound on the perturbation induced by such Gaussian coherent noise based on the Hessian matrix of the circuit, and theoretically and numerically show that hardware-efficient ansätze may be particularly resilient against this type of error due to small second derivatives [36]. Finally, we analyse the performance of the above algorithms under the action of incoherent errors coming from the unavoidable interaction of the qubits with the environment which we have no control over. To study this type of noise, we start by investigating the effect of single-qubit depolarization channels. In addition, we consider a custom noise model that combines various types of errors present on hardware, and study the effect of this noise model with error probabilities that are present in currently available superconducting quantum hardware. Our results show that both policy gradient methods and Q-learning exhibit a robustness to noise that may enable successfully running them on near-term devices. This motivates further study in the quest to find a real-world problem of interest where a quantum advantage for variational RL could be possible.

2 Reinforcement learning

In this section, we will provide a brief introduction to RL that contains the basics necessary to understand this work. For a more in-depth introduction to the topic we refer the reader to [37].

In RL, an agent learns to perform a specific task by trial and error through interacting with an environment. In contrast to supervised learning, this means that there is no necessity for a preexisting training dataset made of pairs of inputs and corresponding correct labels. Instead, the learning task is specified in terms of an *environment* and a *reward function*. The environment is defined in terms of its state space \mathcal{S} and its action space \mathcal{A} , as well as a transition function $P_{ss'}^a = P(s'|s, a)$ that specifies the probability of transitioning to state s' , given that the previous state is s and action a is taken. The agent can use the actions $a \in \mathcal{A}$ to move across states $s \in \mathcal{S}$ of the environment, and receives a reward r that informs it about the quality of the chosen action. The agent chooses its actions based on a policy $\pi(a|s)$ which specifies the probability of taking actions given states, and its goal is to maximize the rewards. This is formally defined as a quantity called the *expected return*, which is the random variable G_t ,

$$G_t = \sum_{k=0}^{\infty} \gamma^k \cdot r_{t+1+k}, \quad (1)$$

where $\gamma \in [0, 1)$ is a discount factor that controls the significance of delayed rewards, t is the current time step and r_t represents the reward at the given time step. Typically we work in *episodic* environments with a fixed time *horizon* H , so that the sum in Equation (1) runs until H instead of infinity. We can then quantify the agent's performance in terms of a *value function*,

$$V_{\pi}(s) = \mathbb{E}_{\pi} \left[\sum_{k=0}^{H-1} \gamma^k \cdot r_{t+1+k} \mid s_t = s \right], \quad (2)$$

which is the expected return when following a given policy π from an initial state s . There are many different approaches to maximize the expected return, and we focus on the two main paradigms used in state-of-the-art RL: value-based and policy gradient methods. We will now introduce both of these in more detail.

2.1 Value-based methods

One approach to maximizing the expected return is to parameterize and train the value function in Equation (2) directly with a function approximator. This function approximator can be implemented for example as a neural network (NN) [38] or a parameterised quantum circuit (PQC) [26, 27, 29]. The value-based method that we focus on in this work is called *Q-learning*. While the value function in Equation (2) is called the *state-value* function as it only depends on the state, in Q-learning we try to approximate the *action-value* function that additionally depends on the action,

$$Q_{\pi}(s, a) = \mathbb{E}_{\pi} \left[\sum_{k=0}^{H-1} \gamma^k \cdot r_{t+1+k} \mid s_t = s, a_t = a \right]. \quad (3)$$

For a parametrized Q-function $Q_\pi(s, a; \theta)$ the goal is then to approximate the optimal Q-function Q^* as closely as possible, where the optimal Q-function is the one that leads to the optimal policy. The actions are chosen such that in each time step, the agent prefers to take the action that has the highest expected return, i.e.,

$$a_t = \operatorname{argmax}_a Q_\pi(s_t, a; \theta). \quad (4)$$

Due to this choice being deterministic, a Q-learning agent may never visit certain states of the environment and therefore not explore the state space sufficiently to find a good policy. In order to facilitate exploration, in practice a so-called ϵ -greedy policy is used, where the agent selects a random action instead of that corresponding to the largest Q-value with probability ϵ . Typically, ϵ is chosen large at the beginning and decreased over the course of training. In each training step, the Q-values are updated as follows,

$$Q_\pi(s_t, a_t; \theta) \leftarrow r_{t+1} + \gamma \max_a Q_\pi(s_{t+1}, a; \theta). \quad (5)$$

In order to train a function approximator like a NN or a PQC, the right-hand side of Equation (5) is used as a label in a supervised learning setting. This means that the function approximator is updated based on its own predictions about the expected return under the current parametrization, in addition to the reward given by the environment. Consequently, the agent needs to learn a moving target, which can lead to instability of training and delayed convergence. Additionally, updates are always based on the latest observed rewards, so the agent can “forget” previously learned behaviour even when it was beneficial.

To stabilize training, two components have been added to the algorithm: a second model to compute the Q-values on the right-hand side of Equation (5), called the *target model*, which is identical to the Q-function approximator but with parameters that are updated with a copy of the Q-function approximator’s parameters only at fixed intervals. This decreases the rate of change in the prediction of the expected return used for parameter updates, and can therefore make learning more stable. Additionally, past interactions with the environment are stored in a memory and then randomly sampled to perform parameter updates to remove temporal correlations between transitions. For more detail on Q-learning with function approximators, also referred to as *deep Q-learning* in classical literature, we refer the reader to the seminal work [38].

2.2 Policy gradient method

As described above, a RL agent chooses its actions based on a policy $\pi(a|s)$, which is the conditional probability distribution of actions given states. To maximize the expected return, the agent needs to find the optimal policy π^* . In policy gradient training, the agent is implemented in form of a parametrized policy π_θ , and the goal of the algorithm is to find the parameters θ^* that produce the optimal policy. The quality of the policy is measured by a quantity $J(\theta)$, that in the fixed-horizon setting is equal to the value function (2),

$$J(\theta) := V_{\pi_\theta}(s). \quad (6)$$

In a gradient-based optimization procedure the parameters are updated according to

$$\boldsymbol{\theta}_{t+1} = \boldsymbol{\theta}_t + \alpha \nabla J(\boldsymbol{\theta}_t), \quad (7)$$

with a learning rate α , i.e., we perform gradient ascent on the parameters to maximize the expected return. The policy gradient theorem [37] then states that the gradient of our performance measure can be written as

$$\begin{aligned} \nabla J(\boldsymbol{\theta}) &= \nabla V_{\pi_{\boldsymbol{\theta}}}(s) \\ &\propto \sum_s \mu(s) \sum_a \nabla \pi_{\boldsymbol{\theta}}(a|s) Q_{\pi}(s, a) \\ &= \mathbb{E}_{\pi_{\boldsymbol{\theta}}} \left[\sum_a \nabla \pi_{\boldsymbol{\theta}}(a|S_t) Q_{\pi}(S_t, a) \right], \end{aligned} \quad (8)$$

where $\mu(s)$ is the on-policy distribution under the current policy, which depends on the time spent in each state, and S_t in the third line of Equation (8) are states sampled under the policy π . Using this, we can now derive the REINFORCE algorithm, that is the basis of policy gradient based training.

Our goal is to perform gradient ascent on the parametrized policy purely from samples generated from said policy through interactions with the environment. The last line of Equation (8) still contains a sum over all actions a , which we can replace by the sample $A_t \sim \pi$ after multiplying and dividing the terms in the sum by $\pi_{\boldsymbol{\theta}}(a|S_t)$,

$$\begin{aligned} \nabla J(\boldsymbol{\theta}) &\propto \mathbb{E}_{\pi_{\boldsymbol{\theta}}} \left[\sum_a \pi_{\boldsymbol{\theta}}(a|S_t) Q_{\pi}(S_t, a) \frac{\nabla \pi_{\boldsymbol{\theta}}(a|S_t)}{\pi_{\boldsymbol{\theta}}(a|S_t)} \right] \\ &= \mathbb{E}_{\pi_{\boldsymbol{\theta}}} \left[Q_{\pi}(S_t, A_t) \frac{\nabla \pi_{\boldsymbol{\theta}}(A_t|S_t)}{\pi_{\boldsymbol{\theta}}(A_t|S_t)} \right] \\ &= \mathbb{E}_{\pi_{\boldsymbol{\theta}}} \left[G_t \frac{\nabla \pi_{\boldsymbol{\theta}}(A_t|S_t)}{\pi_{\boldsymbol{\theta}}(A_t|S_t)} \right], \end{aligned} \quad (9)$$

where G_t is the expected return from Equation (1). Now, by using the fact that $\nabla \log x = \frac{\nabla x}{x}$, we can write

$$\mathbb{E}_{\pi_{\boldsymbol{\theta}}} \left[G_t \frac{\nabla \pi_{\boldsymbol{\theta}}(A_t|S_t)}{\pi_{\boldsymbol{\theta}}(A_t|S_t)} \right] = \mathbb{E}_{\pi_{\boldsymbol{\theta}}} [G_t \cdot \nabla \log \pi_{\boldsymbol{\theta}}(A_t|S_t)]. \quad (10)$$

This equation allows us to estimate the gradient of $J(\boldsymbol{\theta})$ by samples from the current policy $\pi_{\boldsymbol{\theta}}$, and leads us to the following parameter update in each iteration of the algorithm,

$$\boldsymbol{\theta} \leftarrow \boldsymbol{\theta} + \alpha \gamma^t \sum_{k=t+1}^T \gamma^{k-t-1} R_k \nabla \log \pi_{\boldsymbol{\theta}}(A_t, S_t), \quad (11)$$

where α is again the learning rate, R_k is the reward, and T is the length of the episode. Quantum versions of policy gradient based learning have been introduced in [28, 32], where the policy is implemented in form of a PQC.

3 Environments and implementation

Our goal is to get insight into the effect of noisy training on quantum RL algorithms. For this, we consider quantum versions of the two main paradigms in RL that have been introduced in previous sections: value-based methods (see Sect. 2.1) and policy gradient methods (see Sect. 2.2). As we are interested in the effect of noisy training on models that have otherwise been proven to work well in the noise-free setting, we study models and environments that have been already investigated in this setting before [28, 29, 39]. In this way, we have evidence that the models and hyperparameters that we choose are suitable for the studied environments, and can focus our efforts on understanding the effect that noise has on the training and performance of these agents. The code that was used to generate the numerical results in this work can be found on Github [40].

3.1 CartPole

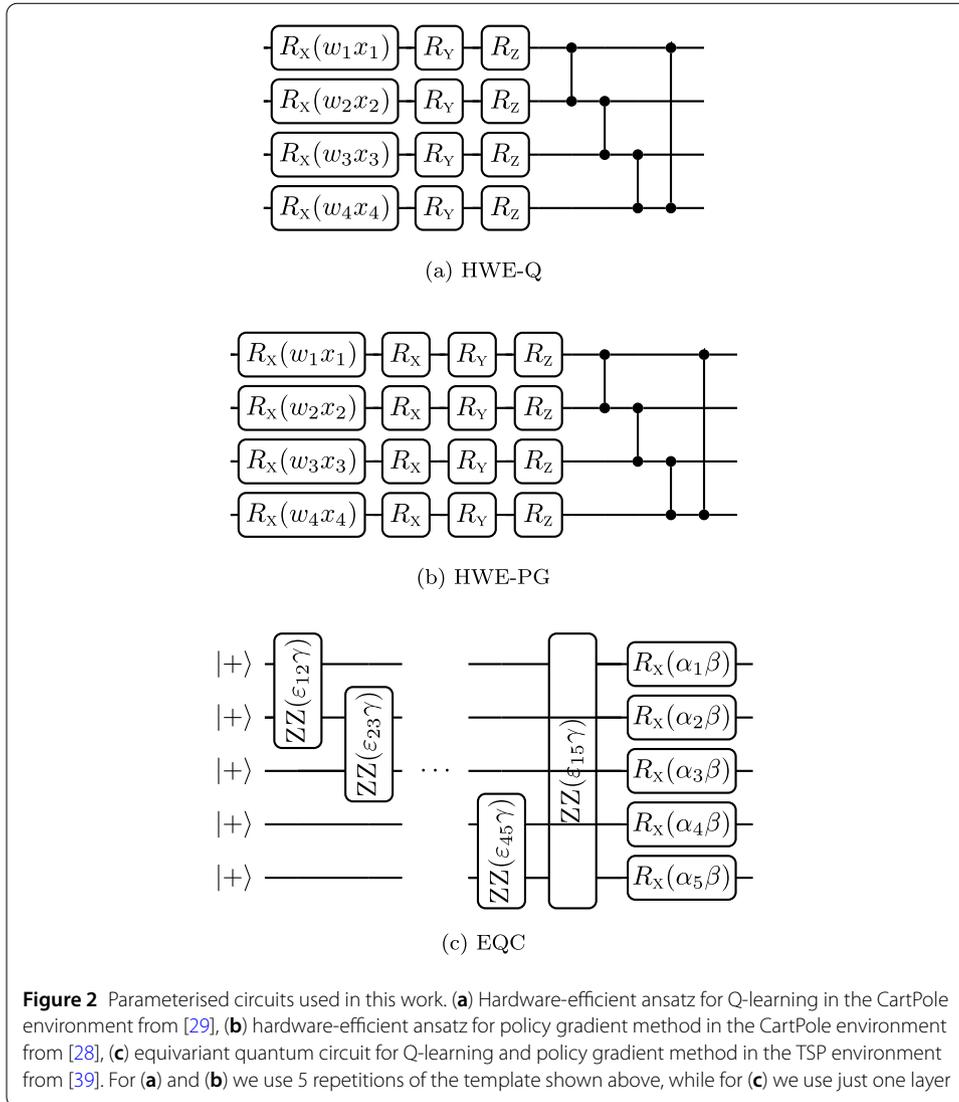
The first environment that we study is a benchmark task from the classical literature and implemented in the OpenAI Gym [41]: the CartPole environment. It has been previously studied in classical and quantum RL literature [27–29, 42]. In this environment, the goal is to learn to balance a pole that is attached to a cart that can move left and right on a frictionless track. The state s of the environment is represented by a four dimensional input vector $s \rightarrow \mathbf{x} = (x_1, x_2, x_3, x_4) \in \mathbb{R}^4$ encoding the position and velocity of the cart, and the velocity and angle of the pole. There are two actions that the agent can perform: moving the cart left or right. The environment is considered as solved when the agent manages to balance the pole for an average of at least 195 time steps for 100 consecutive episodes. We implement noisy training for the CartPole environment using the policy gradient approach introduced in [28] and the Q-learning approach introduced in [29].

The circuit used for Q-learning in [29] consists of five layers of a hardware-efficient ansatz [43], where each circuit layer consists of one parametrized rotation around the x -axis per qubit that is used to encode the input states \mathbf{x} , and additional parametrized y - and z -rotations on each qubit that contain the free parameters to be trained (see Fig. 2(a)). Furthermore, additional trainable parameters multiplying each input feature are used to increase the expressivity of the reuploading quantum circuit [44, 45]. Each layer also has a final layer of CZ-gates arranged in a circular topology. The observable for taking the action “left” is $O_L = Z_1 Z_2$, where Z_1 and Z_2 are Pauli-Z operators acting on the first and second qubit, respectively. Similarly, action “right” is associated to the observable $O_R = Z_3 Z_4$, defined on the third and fourth qubit. In order to facilitate the function approximation of the optimal Q-function, which has a range of output values beyond that of $Z_i Z_j$ operators, each expectation value is further multiplied with an additional trainable weight, such that the final Q-value for action “left” is

$$Q(s, L) = \frac{\langle O_L \rangle_{s, \theta} + 1}{2} w_L \quad (12)$$

$$= \frac{\langle \mathbf{0} | U_\theta^\dagger(s) O_L U_\theta(s) | \mathbf{0} \rangle + 1}{2} w_L, \quad (13)$$

where $U_\theta(s)$ represents the unitary of the parameterised circuit depending on the trainable parameters θ and the input state s , and w_L is the trainable weight correspond-



ing to observable O_L . The Q-value for the action “right” is defined in a similar manner.

For the policy gradient method, we follow the implementation used in [28] and made available at [46], which uses five layers of the same hardware-efficient ansatz as described for Q-learning above, except that each layer has an additional trainable rotation around the x -axis on each qubit (see Fig. 2(b)), and the actions observables are defined as $O_L = Z_1 Z_2 Z_3 Z_4$ and $O_R = \mathbb{I} - O_L$. As before, input features are multiplied with an additional trainable parameter each. Since the policy is a probability distribution, a final SOFTMAX layer is used to map the expectation values $\langle O_a \rangle_{s,\theta} \in [-1, 1]$ to the appropriate range $[0, 1]$, and so probabilities for each action eventually become

$$\pi_\theta(a|s) = \frac{e^{\beta \langle O_a \rangle_{s,\theta}}}{\sum_{a'} e^{\beta \langle O_{a'} \rangle_{s,\theta}}}, \tag{14}$$

where $\beta \in \mathbb{R}$ is also a trainable parameter.

3.2 Traveling salesperson problem

The second environment that we study is more complex and requires introducing the field of neural combinatorial optimization (NCO). NCO is an alternative to the hand-crafted heuristics used in combinatorial optimization, where instead a machine learning model is trained to solve instances of a given combinatorial optimization problem [47]. In the case of RL-based NCO, the optimization problem is defined in terms of an environment and the quality of the solution is measured by the reward function. In this work, we study a quantum NCO approach that learns to solve instances of the Traveling Salesperson Problem (TSP), as introduced in [39]. In TSP one is presented with a list of cities in form of a weighted graph, and the goal is to find the tour of minimal length that visits each city in this list exactly once.

In this environment one episode consists in solving one instance, where the agent selects the cities in the tour in a step-wise fashion. States in this environment are instances of the TSP, in addition to the partial tour at the current time step. The actions are defined in terms of the cities, where in each time step the agent can select one of the cities that is not yet in the tour. The reward is the negative difference in length between the tour at the previous time step and the tour after adding the latest city, as we want to minimize the length of the tour while RL agents try to maximize the expected reward. We evaluate the quality of the tours proposed by the agents in terms of the approximation ratio

$$\frac{c(T)}{c(T^*)},$$

where $c(T)$ is the length of the tour T proposed by the agent, and $c(T^*)$ is the length of the optimal tour T^* . The stopping criterion for this environment is an average approximation ratio of at least 1.05 over the past 100 episodes.

To implement a quantum agent for this environment, we follow [39], where the information of the TSP graph instance is directly encoded into a PQC and each graph node corresponds to one qubit. Each layer in this ansatz consists of one rotation around the x -axis parametrized by $\alpha_i\beta_i$, where $\alpha_i \in \{0, \pi\}$ represents whether city i is already in the current tour ($\alpha_i = 0$), or still available for selection ($\alpha_i = \pi$), and β_i is a trainable parameter that is shared across all single-qubit gates in layer l . The graph's edges in each layer are represented by a ZZ-gate parametrized by $\varepsilon_{ij}\gamma_l$, where ε_{ij} is the weight of edge connecting nodes i and j , and γ_l is a trainable parameter that is shared across all two-qubit gates in layer l . Such ansatz is shown in Fig. 2(c).

In the case of Q-learning, the observables are ZZ-operators that correspond to the edges in the graph, i.e., Z_iZ_j is measured for edge ij . For policy gradient agents the observables are the same, but as the policy has to be a probability distribution we again use a final SOFTMAX layer with a trainable inverse temperature β , as in Equation (14). The authors of [28] have shown that using this type of final layer can be highly beneficial for policy gradient training, compared to only using the probability distribution resulting from the quantum state directly. This is due to the fact that the trainable inverse temperature enables the agent to tune its level of exploration of the state space. As the optimal solutions to TSP instances are deterministic, it is favourable in this environment to have a tunable inverse temperature that allows exploration of the large state space early in training, as well as close-to-deterministic decisions towards the end.

4 Shot noise

We start our studies with the type of noise that is arguably the simplest to characterize: noise induced by statistical errors that result from the probabilistic nature of quantum measurements. For each circuit evaluation, be it for action selection of the RL agent or for computing parameter updates via the parameter shift rule, we take a fixed number of measurements M and compute the resulting expectation value. The precision of this expectation value depends on M and scales like $\epsilon \sim 1/\sqrt{M}$.

Variational algorithms often require a very large number of measurements to be executed, and this problem is exacerbated in QML tasks that typically involve separate circuit evaluations for all training data points. For this reason, it is not only important to understand the effect of shot noise on the trainability and performance of QML models, but it is also desirable to develop methods that lead to a smaller shot footprint than simply assigning a fixed number of shots to each circuit evaluation. Depending on knowledge of the algorithm itself, it can be possible to make an informed decision on the number of shots that suffice in each step. In this section, we develop such a method specifically for Q-learning that is a natural extension to the original algorithm.

4.1 Reducing the number of shots in a Q-learning algorithm

As described in Sect. 2.1, a Q-learning agent selects actions based on the following rule (see Equation (4))

$$a_t = \operatorname{argmax}_a Q_\pi(s_t, a; \theta),$$

that is, it chooses actions according to the largest Q-value.¹ Now, consider a quantum agent that only has access to noisy estimates of the Q-values $\tilde{Q}(s_t, a_t; \theta)$ resulting from the statistical uncertainty of a measurement process involving a finite number of shots M . If the sample size is large enough $M \gg 1$, then by the central limit theorem each noisy Q-value can be described as a random variable

$$\tilde{Q}(s_t, a_t; \theta) = Q(s_t, a_t; \theta) + \epsilon, \quad (15)$$

where $Q(s_t, a_t; \theta)$ is the true noise-free value, and ϵ is a random variable sampled from a Gaussian distribution centered in zero $\mu_\epsilon = \mathbb{E}[\epsilon] = 0$, and with standard deviation inversely proportional to the square root of the number of measurement shots $\sigma_\epsilon = \operatorname{Std}[\epsilon] \sim 1/\sqrt{M}$. Since actions are selected through an argmax function, the perturbation ϵ will not affect the action selection process as long as the order between the largest and the remaining Q-values remains unchanged. Then, one may ask: is there a minimal number of shots that suffice to reliably distinguish the largest Q-value Q_{max} and the second-largest Q-value Q_2 ?

When the observables associated to the actions are non-commuting, they have to be estimated independently from each other, and one has the freedom of choosing how to allocate the measurement shots among the observables of interest, possibly in a clever

¹In the ϵ -greedy policy (see Sect. 2.1) we consider here, the agent picks either the action corresponding to the argmax Q-value, or a random action. As no circuit evaluation is required to pick a random action, we only consider the steps with actual action selection by the agent in this section.

way. In our case, the goal is to estimate which of the observables has the highest Q-value while trying to be shot-frugal, and this task can be related to the theory of multi-armed bandits [48]. The multi-armed bandit is a RL problem in which an agent can allocate only a limited amount of resources between a number of choices, e.g., a number of arms on a bandit machine, and is asked to determine which of these choices leads to the highest expected reward. There exists a trade-off between exploration (i.e., trying the different arms) and exploitation (always choosing the arm that appears best according to the current knowledge), and the upper confidence bound (UCB) [49, 50] algorithm shows how to use statistical confidence bounds to allocate exploratory resources. The UCB algorithm could be used in the scenario described above where a number of non-commuting observables have to be estimated, and we want to find the optimal strategy to allocate a fixed budget of measurement shots to the task of identifying the largest Q-value.

However, in the specific implementations of QRL agents based on recent literature that we study in this work [28, 29, 39], only commuting observables are used, hence it is not necessary to apply the UCB procedure to determine which one should be measured more often. Nonetheless, inspired by the UCB algorithm, we can still define a rather general simple heuristic that can be used to reduce the overall number of shots required to train the Q-learning models as those studied in this work. The idea is to use the knowledge about the scaling of the estimation error with respect to the number of measurements (see Equation (15)), to determine with confidence whether we have taken enough shots to determine the maximum Q-value.

The procedure goes as follows. First, we take a small number of initial measurements m_{init} , for example $m_{\text{init}} = 100$, of all observables to compute the estimates $\tilde{Q}_{m_{\text{init}}}(s_t, a)$, $\forall a \in \mathcal{A}$. Based on these values, we compute the absolute difference between the largest and the second largest Q-values. If this difference is larger than twice the estimation error $\epsilon = 2/\sqrt{m_{\text{init}}}$ (as both of the Q-values are noisy), we have found the largest Q-value with high confidence and we stop here. On the other hand, if the difference is smaller, we increment the sample size with additional m_{inc} measurements each, and recompute the estimated Q-values with the $m_{\text{inc}} + m_{\text{init}}$ shots. We again compute the absolute difference of the two largest Q-values and determine whether the number of measurements suffices based on the error $\epsilon = 2/\sqrt{m_{\text{init}} + m_{\text{inc}}}$. This measure-and-compare scheme is performed until either the two largest Q-values can be distinguished with high confidence, or a fixed shot budget m_{max} is reached.

In Algorithm 1 we provide a description of this procedure, where for the sake of simplicity we describe the case where there are only two possible actions, and we therefore only have to find the larger of two Q-values. However, the scheme can be used for an arbitrary number of Q-values, as it is only important to distinguish between the highest and the second-highest Q-value with high confidence. The algorithm takes as input the number of initial measurements m_{init} , the number of additional measurements in every step m_{inc} , and the maximum number of measurements that are allowed in one run of the shot-allocation algorithm (i.e., finding the largest Q-value) m_{max} . The output is the number of measurements m_{est} that are sufficient to find the argmax Q-value with high confidence based on the rules above. The values $\langle O_{a_i} \rangle_{m_{\text{est}}}$ are the expectation values of observables O_{a_i} corresponding to action a_i , estimated with m_{est} shots. Note that the proposed scheme works both for commuting or non-commuting observables, where in the former case one

Algorithm 1: Algorithm to reduce the number of measurements in Q-learning

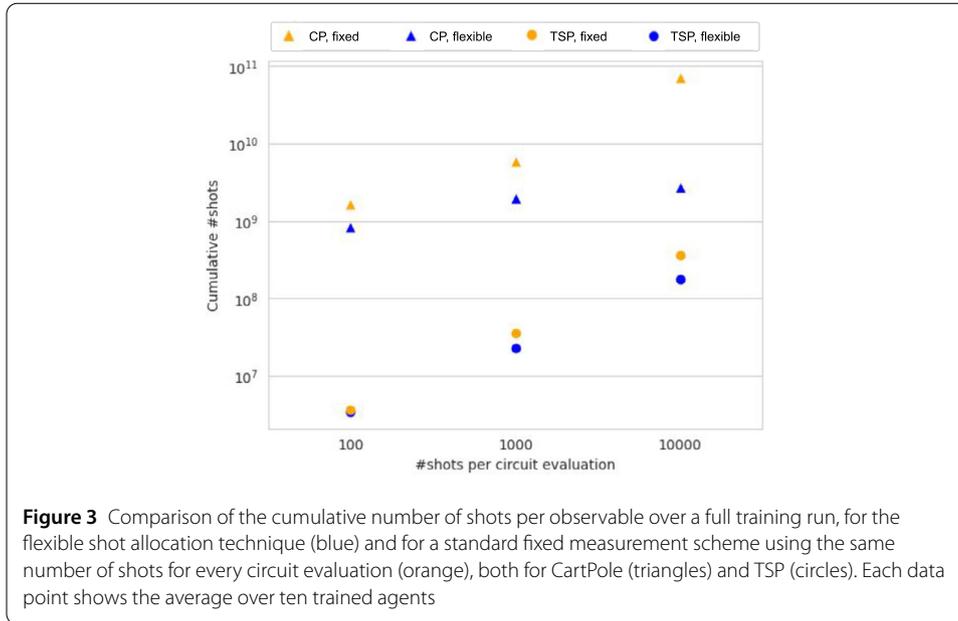
Input : $m_{\text{init}}, m_{\text{inc}}, m_{\text{max}}$
Output: m_{est}
 $m_{\text{est}} \leftarrow m_{\text{init}}$
while $m_{\text{est}} < m_{\text{max}}$ **do**
 $\tilde{Q}(s_t, a_1) = \langle O_{a_1} \rangle_{m_{\text{est}}}$;
 $\tilde{Q}(s_t, a_2) = \langle O_{a_2} \rangle_{m_{\text{est}}}$;
 $\Delta \tilde{Q} = |\tilde{Q}(s_t, a_1) - \tilde{Q}(s_t, a_2)|$;
 if $\Delta \tilde{Q} < 2/\sqrt{m_{\text{est}}}$ **then**
 $m_{\text{est}} \leftarrow m_{\text{est}} + m_{\text{inc}}$;
 else
 return $\min(m_{\text{est}}, m_{\text{max}})$;
 end
end
return $\min(m_{\text{est}}, m_{\text{max}})$;

can spare shots by computing the observables from the same set of measurement outcomes. Moreover, note that we ignore the coefficients in the statistics of the Q-values coming from Equation (12), when considering the measurement stopping criterion. This choice has no impact on the effectiveness of the proposed method, as it is always found to be very well performing in the presented form.

While this algorithm can clearly determine the optimal number of shots in the action selection process in a methodical manner, one should check that this will not introduce errors in the remaining parts of the variational Q-learning model, i.e., during the parameter update step. Recall that each parameter update of the model is computed based on the output of the model itself (see Equation (5))

$$Q_{\pi}(s_t, a_t; \theta) \leftarrow r_{t+1} + \gamma \max_a Q_{\pi}(s_{t+1}, a; \theta),$$

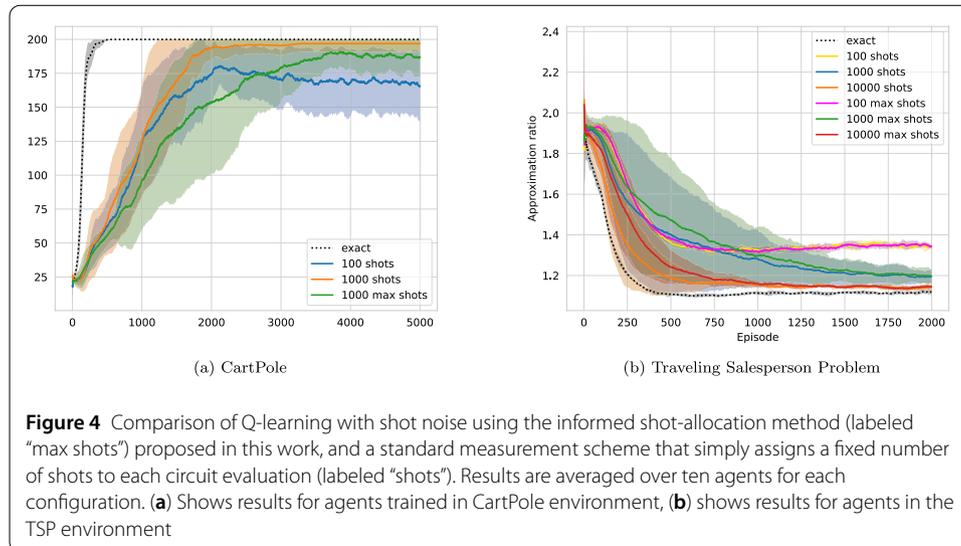
which means that in the parameter update step we do not need to perform action selection, but instead care about the actual Q-values in order to compute the loss function. The question is now to what precision we need to approximate the Q-values in order to learn a good Q-function. Technically, even the noise-free Q-function is only an approximation of the true Q-function, which is the whole point of doing Q-learning with function approximators. This suggests that there is some leeway to make even the approximate function itself an approximation by taking only as many measurements as are necessary to find the argmax Q-value with high confidence. Indeed, it has been shown in [29] that even the Q-functions of agents that successfully solve an environment can produce Q-values that are far from the optimal Q-values, and that learning the correct order of Q-values is more important in this setting than approximating the optimal Q-value as precisely as possible. Consequently, when we compute the Q-values that are used to perform parameter updates, we use the same algorithm as that in Algorithm 1 to determine the number of measurements to take.



4.2 Numerical results

We now numerically compare the performance of agents in the CartPole and TSP environments in settings where a fixed number of shots is used in each circuit evaluation, and where the number of shots in each step is determined by the algorithm we introduced in Sect. 4.1. To give an overview of the number of shots used in one training run under varying hyperparameter settings, we show the average cumulative number of shots for different settings in Fig. 3. For the CartPole environment (triangles), the number of cumulative shots grows quickly with the number of shots in each step in the fixed setting (orange). This is not true for the flexible shot allocation technique (blue), where for values of $m_{\max} \in \{100, 1000, 10,000\}$ the cumulative number of shots is relatively similar. As we see in Fig. 4 a), a low number of shots such as 1000 is already sufficient to achieve close to optimal performance in the CartPole environment. Therefore, we focus on comparing settings with 100 and 1000 (maximum) shots per circuit evaluation in that figure. Comparing the cumulative number of shots for $m_{\text{fixed}} = 100$ and $m_{\max} = 1000$ in Fig. 3, we see that these two configurations use almost the same number of measurements overall. Still, the final performance of the agents trained with the flexible shot allocation technique is almost optimal, while those trained with a fixed number of shots in each circuit evaluation are below a final score of 175 on average. However, as we allow agents to use even less than 100 shots per evaluation with the flexible allocation method of Algorithm 1, performance starts to degrade, so at least 100 shots are required in this setting. To not clutter the figure we show the results for agents that use fewer than 100 shots per circuit evaluation in Fig. 15 in the Appendix.

In the TSP environment, each step in an episode constitutes of a constant and (compared to CartPole) relatively low number of circuit evaluations. We still see that the higher the setting for the (maximum) number of shots is, the bigger the gap in average cumulative number of shots becomes. For agents trained in the TSP environment, shown in Fig. 4 b), the final performance remains unchanged by the additional noise introduced by the flexible shot allocation technique, and agents reach the same accuracy of those trained with



a corresponding but fixed number of shots per circuit evaluation. The only difference between the two approaches is that the agents using the flexible shot allocation method take slightly longer to converge in some cases. Independently from the estimation method used (flexible or fixed), it is clear from Fig. 4 that it is the number of shots available that plays the major role in determining the performance of the noisy agents, as measured by the proximity to the average approximation ratios reached in the noise-free scenario, namely when agents have access to exact the expectation values ($M \rightarrow \infty$). In this environment, there is a trade-off between delayed convergence due to less precision in the approximation of the Q-function, and using a higher number of shots to arrive at the same final performance.

To summarize, we have seen that Q-learning models can be successfully trained even in the presence of statistical noise introduced by a measurement processes carried out with a limited number of shots. In addition, by leveraging the specifics of the Q-learning algorithm, we introduced an easy-to-implement and effective method that can be used to reduce the number of shots needed to train variational Q-learning agents. How many shots one can save during training with this method depends on the agents’ resilience to shot noise, as well as the specific characteristics of the environment. In the CartPole environment, where one bad decision does not lead to immediate failure, the additional noise introduced by estimating expectation values with a low number of measurements and approximating an imprecise Q-function does not affect performance severely. In the TSP environment on the other hand, where one bad choice of the next city in the tour can lead to a much longer path, we observe that the number of measurements has to be relatively high to get close to optimal performance. However, even in this setting we can achieve a reduction in the overall number of measurements by taking an informed approach at when to measure an observable more often.

5 Coherent noise

In this section, we turn our attention to coherent noise, that is, errors that preserve the unitary evolution of the quantum circuit but still change its output [51]. In our analysis, we model coherent noise as an *over-* or *under-rotation* of the parametrized gates, by adding a random Gaussian perturbation to the variational parameters in the considered circuits.

This type of error could occur in real quantum devices as a drift in the parameters for example due to an imperfect control of the system or a miscalibration of the hardware, and it is therefore an important component of the overall picture of an imperfect quantum device. Specifically, we assume that the perturbation remains unchanged during the estimation of a given observable, i.e. it does not change considerably between repeated measurements on the same experimental setup. However the perturbation amount changes whenever the experiment is changed, for example due to measuring a different observable, or using the circuit with a different set of parameters.

Gaussian coherent noise is also an interesting model because it lends itself very well to theoretical analysis, and one can estimate the effect of such an error on the output of a parameterised quantum circuit. In the following, we first proceed with an analytical treatment of the error introduced by Gaussian perturbations on variational circuits, and then proceed with the numerical results for the two environments considered in this work.

5.1 Effect of Gaussian coherent noise on circuit output

Consider a general parametrized quantum circuit acting on a system of n qubits, with unitary $U(\boldsymbol{\theta}) \in \mathbb{C}^{2^n} \times \mathbb{C}^{2^n}$ and parameter vector $\boldsymbol{\theta} = (\theta_1, \dots, \theta_M) \in \mathbb{R}^M$. Let O be an observable and $\rho = |0\rangle\langle 0|$ the initial state of the quantum system, the outcome of the variational circuit is the expectation value

$$f(\boldsymbol{\theta}) = \langle O \rangle_{\boldsymbol{\theta}} = \text{Tr}[OU(\boldsymbol{\theta})\rho U^\dagger(\boldsymbol{\theta})]. \quad (16)$$

Suppose that the parameters are affected by a noise process that adds a perturbation

$$\boldsymbol{\theta} \rightarrow \boldsymbol{\theta} + \delta\boldsymbol{\theta}, \quad (17)$$

where $\delta\boldsymbol{\theta} = (\delta\theta_1, \dots, \delta\theta_M) \in \mathbb{R}^M$ are *i.i.d.* according to a Gaussian distribution $\mathcal{N}(\mu, \sigma)$ with zero mean $\mu = 0$ and equal variance σ^2 , namely

$$\begin{aligned} \delta\theta_i &\sim \mathcal{N}(0, \sigma^2), \\ \mathbb{E}[\delta\theta_i] &= 0, \quad \forall i \in \{1, \dots, M\}, \\ \mathbb{E}[\delta\theta_i \delta\theta_j] &= \sigma^2 \delta_{ij}, \end{aligned} \quad (18)$$

As discussed earlier, in our analysis in this section and in the numerical simulations in Sect. 5.3.1, we assume that the perturbed parameters remain the same during the evaluation of a single expectation value. In a real experiment on quantum hardware, this would mean that for all measurements used to estimate the expectation value, the perturbations stay at least approximately unchanged. Of course, without this assumption, the resulting noise model could not be considered unitary, and one may then resort to a noise channel formulation of Gaussian noise as proposed in [4, 35]. Hence, in the following we restrict our attention to the setting described above.

The effect of Gaussian noise on the circuit can be evaluated by Taylor expanding the circuit around the unperturbed parameters $\boldsymbol{\theta}$. For ease of explanation, we hereby report only the main ideas and results, and we refer to Appendix B for a complete and detailed derivation of all the calculations performed in this section.

Let $f(\boldsymbol{\theta} + \delta\boldsymbol{\theta})$ be the function evaluated on the perturbed parameters, its Taylor expansion up to fourth-order reads

$$\begin{aligned} f(\boldsymbol{\theta} + \delta\boldsymbol{\theta}) &\approx f(\boldsymbol{\theta}) + \sum_{i=1}^M \frac{\partial f(\boldsymbol{\theta})}{\partial \theta_i} \delta\theta_i + \frac{1}{2} \sum_{i,j=1}^M \frac{\partial^2 f(\boldsymbol{\theta})}{\partial \theta_i \partial \theta_j} \delta\theta_i \delta\theta_j \\ &+ \frac{1}{3!} \sum_{i,j,k=1}^M \frac{\partial^3 f(\boldsymbol{\theta})}{\partial \theta_i \partial \theta_j \partial \theta_k} \delta\theta_i \delta\theta_j \delta\theta_k + \mathcal{O}(\delta\theta^4). \end{aligned} \quad (19)$$

With this expression one can evaluate the expected value of the noisy function $\mathbb{E}[f(\boldsymbol{\theta} + \delta\boldsymbol{\theta})]$ over the distribution of the Gaussian perturbations, $\mathbb{E}(\cdot) = \mathbb{E}_{\delta\theta_i \sim \mathcal{N}(0, \sigma^2)}(\cdot)$. Since every odd moment of a Gaussian distribution vanishes, using relations (18) in the expansion (19) one obtains

$$\begin{aligned} \mathbb{E}[f(\boldsymbol{\theta} + \delta\boldsymbol{\theta})] &\approx f(\boldsymbol{\theta}) + \frac{1}{2} \sum_{ij} \frac{\partial^2 f(\boldsymbol{\theta})}{\partial \theta_i \partial \theta_j} \mathbb{E}[\delta\theta_i \delta\theta_j] \\ &\approx f(\boldsymbol{\theta}) + \frac{1}{2} \sigma^2 \sum_{ij} \frac{\partial^2 f(\boldsymbol{\theta})}{\partial \theta_i \partial \theta_j} \delta_{ij} \\ &\approx f(\boldsymbol{\theta}) + \frac{1}{2} \sigma^2 \text{Tr}[H(\boldsymbol{\theta})] + \mathcal{O}(\sigma^4), \end{aligned} \quad (20)$$

where $\text{Tr}[H(\boldsymbol{\theta})]$ denotes the trace of the Hessian matrix

$$H_{ij}(\boldsymbol{\theta}) = \frac{\partial^2 f(\boldsymbol{\theta})}{\partial \theta_i \partial \theta_j}, \quad i, j = 1, \dots, M. \quad (21)$$

Thus, the first non-vanishing correction term caused by the noise is proportional to the noise variance σ^2 , and the Hessian of the parametrized quantum circuit, which conveys geometric information about the curvature of the function landscape around the unperturbed point $\boldsymbol{\theta}$.

Higher-order terms in the expansion can be evaluated in a similar way, specifically making use of so-called *Wick's relations* for multivariate normal distributions as shown in Appendix B. If all the derivatives of the function $f(\boldsymbol{\theta})$ are bounded, as it is the case for parametrized quantum circuits, then it is possible to derive an upper bound on the error induced by the perturbations which only depends on the noise strength σ^2 and the total number of parameters M , as we show in the following.

Using the parameter shift rule [52, 53], one can show that any derivative of a parametrized quantum circuit can be expressed as a linear combination of circuit outcomes evaluated at specific points in parameter space [35, 36]. Let $\boldsymbol{\alpha} = (\alpha_1, \dots, \alpha_M) \in \mathbb{N}^M$ be a multi index keeping track of the order of partial derivatives, define the derivative operator

$$\partial^{\boldsymbol{\alpha}} := \frac{\partial^{|\boldsymbol{\alpha}|}}{\partial \theta_1^{\alpha_1} \dots \partial \theta_M^{\alpha_M}}, \quad (22)$$

where $|\boldsymbol{\alpha}| := \sum_{i=1}^M \alpha_i$. By nested applications of the parameter shift rule, one can show that

$$\partial^{\boldsymbol{\alpha}} f(\boldsymbol{\theta}) = \frac{1}{2^{|\boldsymbol{\alpha}|}} \sum_{i=1}^{2^{|\boldsymbol{\alpha}|}} s_m f(\boldsymbol{\theta}_m), \quad (23)$$

where $s_m \in \{\pm 1\}$ are signs, and θ_m are parameters obtained shifting the parameter vector θ along different directions. Now, since the measurement outcome of every circuit is bounded by the maximum absolute eigenvalue of the observable, i.e. $|f(\theta)| \leq \|O\|_\infty$, consequently it also holds that $|\partial^\alpha f(\theta)| \leq \|O\|_\infty$ (see Appendix B). Note that we only consider bounded observables here, like the Pauli operators commonly used in variational RL algorithms [26–29].

Since all the derivatives of the function are bounded, it is possible to bound every term in the Taylor series and then compute an upper bound to the error caused by the perturbation. In fact, defining the absolute (average) error caused by the noise as

$$\varepsilon_\theta := \left| \mathbb{E}[f(\theta + \delta\theta)] - f(\theta) \right|, \quad (24)$$

one can prove that this is upper bounded by (see Appendix B)

$$\varepsilon_\theta \leq \|O\|_\infty (e^{\sigma^2 M/2} - 1). \quad (25)$$

Note that since $\varepsilon_\theta \leq 2\|O\|_\infty$ is always true, the bound is informative only as long as $e^{\sigma^2 M/2} - 1 < 2$.

This expression only depends on the noise strength σ^2 , the total number of noisy parameters M , and the operator norm of the observable $\|O\|_\infty$, and it can be used to estimate a *sufficient* condition on the noise strength to guarantee a desired error threshold ε_θ . Rearranging Equation (25), a sufficient condition to have error ε_θ not larger than ϵ , is to have Gaussian perturbations satisfying

$$\sigma \leq \sqrt{\frac{2}{M}} \log\left(1 + \frac{\epsilon}{\|O\|_\infty}\right). \quad (26)$$

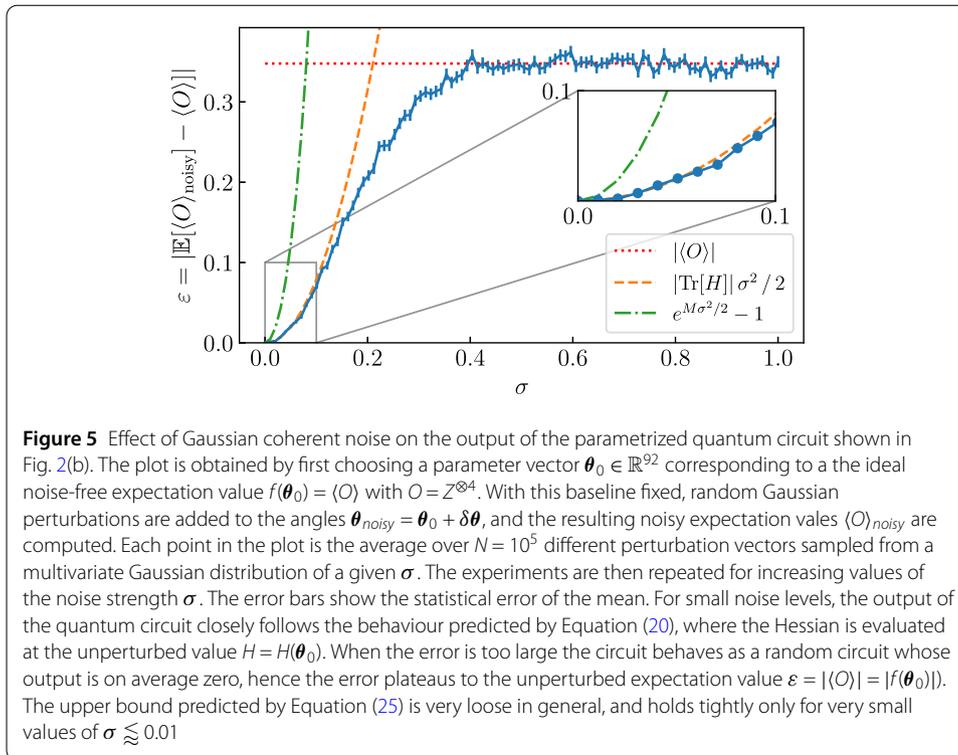
As the allowable error is small $\epsilon \ll 1$, by approximating the logarithm $\log(1+x) \approx x$, one derives that the perturbations must follow the scaling

$$\sigma \in \mathcal{O}\left(\frac{\epsilon}{M^{1/2}\|O\|_\infty}\right). \quad (27)$$

Note that a similar scaling law was recently derived also in [35], though via a slightly different method based on the moment generating function of the probability distribution characterising the perturbations.

To provide an example, assume one is willing to tolerate an error of $\epsilon = 10\%$, that $\|O\|_\infty = 1$ as for measuring a Pauli operator and that the PQC consists of $M = 100$ noisy parametrized gates, then one can be sure of such accuracy if $\sigma \sim 0.1/\sqrt{100} = 0.01$. However, we stress again that the scaling Equation (26) is only a *sufficient* but not necessary condition for achieving an error ϵ . In fact, apart from the requirement of bounded derivatives, Equation (26) is agnostic with respect to the specifics of the function, and such bound can be quite loose in real instances where a much larger noise level still causes a small error, as shown in Fig. 5.

In Fig. 5, we report simulation results obtained by simulating the parametrized ansatz depicted in Fig. 2(b) subject to Gaussian coherent noise of increasing strength. It is clear that the output of the circuit closely follows the approximation of Equation (20) given by



the Hessian even at moderately large value of the noise $\sigma \lesssim 0.15$. When the noise is too strong ($\sigma > 0.2$), the circuit becomes essentially random, and the average expectation value when measuring a Pauli operator is zero. This is a consequence of PQCs often behaving like unitary designs upon random initialization of the parameters [54, 55], a fact which we discuss in detail in Sect. 5.2. At last, as discussed earlier, while the upper bound (25) holds, it is indeed very loose and only holds tightly at small $\sigma \lesssim 0.01$.

We now proceed discussing why hardware-efficient parametrized quantum circuits can be resilient to Gaussian coherent noise. Roughly, this is because such circuits are found to behave like random unitaries upon random assignment of the parameters, which implies that the derivatives of such circuits tend to vanish as the system size grows large [36].

5.2 Resilience of hardware-efficient ansatzes to Gaussian coherent noise

The previous analysis showed that Gaussian perturbations induce an error depending on the Hessian of the circuit (see Equation (20)), so that up to fourth order in the perturbation it holds that

$$\mathbb{E}[f(\theta + \delta\theta)] \approx f(\theta) + \frac{1}{2}\sigma^2 \text{Tr}[H(\theta)]. \tag{28}$$

This equation tells us that if the optimization landscape is flat or close to being flat, then the Hessian is small, and so the perturbation will have little effect on the output of the circuit. On the contrary, in the presence of a very curved landscape, noise will have a great impact and the output of the circuit may change sensibly. It is known that the curvature of the optimization landscape produced by a PQC is closely related to the barren plateau phenomenon [54–56], where the variance of the first and second derivative vanishes exponentially in the number of qubits and layers in a random circuit. Additionally,

the hardware-efficient ansatz we use for some of the environments in this work is known to suffer from barren plateaus when the system size is large. As the curvature of the optimization landscape of these types of circuits is very flat, it can also be expected that the type of noise induced by the Gaussian perturbations on parameters that we study in this work should not affect circuits that generally produce small first and second order derivatives. While circuits that are in the barren plateau regime are obviously undesirable as they quickly become untrainable, one can consider circuits of the size such that the variance in gradients is relatively small, but the circuit has not yet converged to an approximate 2-design, as shown in [54]. We make this statement more formal in the following.

We can use standard results on averages of unitary designs [57, 58] to characterize the Hessian of hardware-efficient circuits, and thus gain insight on their performance under Gaussian noise. We report the main results of our analysis here, full derivations can be found in Appendix B.2. In the following, we suppose that sampling a random value of the parameter vector θ in the parametrized circuit $U(\theta)$, is equivalent to sampling a unitary from a unitary 2-design, defined as a set of unitary matrices that match the Haar distribution up to the second moment. Also, we consider observables O being Pauli strings, so that $\text{Tr}[O] = 0$ and $\text{Tr}[O^2] = 2^n$. In order to distinguish from the previous notation where averages were computed over the Gaussian distribution of the perturbations, we use $\mathbb{E}_U[\cdot]$ and $\text{Var}_U[\cdot]$ to denote average values and variances evaluated over the random unitaries.

Then, under reasonable and usual assumptions on parts of the parametrized quantum circuit being 2-designs, it is possible to show that the diagonal elements of the Hessian $H_{ii} = \partial^2 f(\theta) / \partial \theta_i^2$ satisfy [36] (see also Appendix B.2 for an explicit derivation)

$$\mathbb{E}_U[H_{ii}] = 0, \quad \text{Var}_U[H_{ii}] \in \mathcal{O}\left(\frac{1}{2^n}\right). \quad (29)$$

That is, in addition to first order derivatives, also second order derivatives of random parameterized quantum circuits are found to be zero on average, and with a variance which is exponentially vanishing.

Starting from the results above, one can calculate the statistics of the trace of the Hessian, for which it holds

$$\mathbb{E}_U[\text{Tr}[H]] = 0, \quad \text{Var}_U[\text{Tr}[H]] \lesssim \frac{M^2}{2^n}. \quad (30)$$

Furthermore, our numerical simulations suggest that the variance of the trace of the Hessian is actually smaller, and is well captured by the following expression

$$\text{Var}_U[\text{Tr}[H]] \approx \frac{M(M+1)}{4(2^n+1)} \approx \frac{1}{4} \frac{M^2}{2^n}, \quad (31)$$

a fact which we justify and discuss in Appendix B.2.2.

In Fig. 6 we report simulation results of evaluating the trace of the Hessian matrix for the circuit shown in Fig. 2(b). The histogram represents the frequency of obtaining a given value of the trace of the Hessian $\text{Tr}[H(\theta)]$ upon random assignments of the parameters. Indeed, there is a very good agreement between the variance obtained via numerical simulations (black solid line), and the one calculated with the approximation (31) (dashed red line).

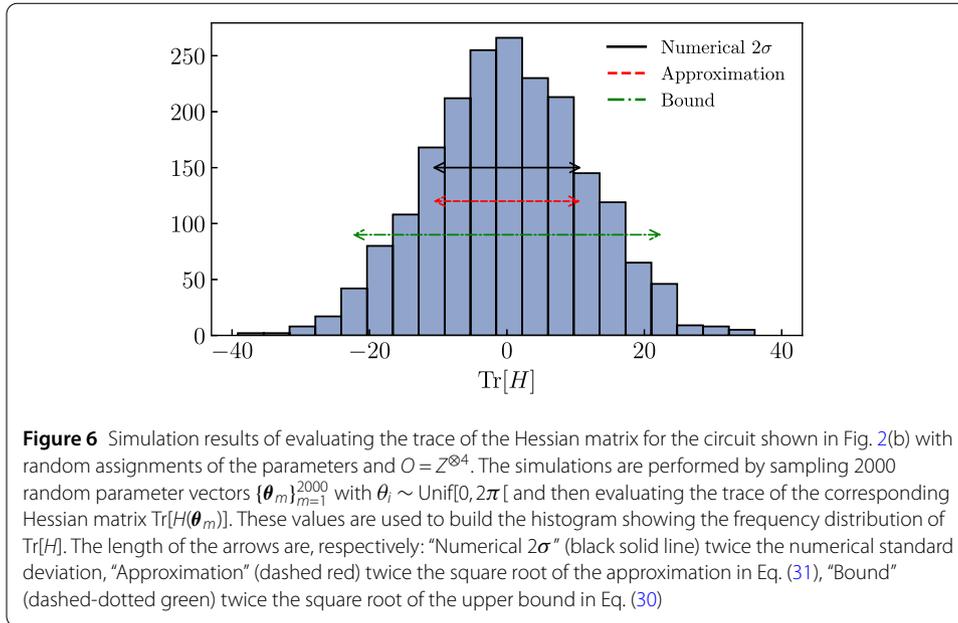


Figure 6 Simulation results of evaluating the trace of the Hessian matrix for the circuit shown in Fig. 2(b) with random assignments of the parameters and $O = Z^{\otimes 4}$. The simulations are performed by sampling 2000 random parameter vectors $\{\theta_m\}_{m=1}^{2000}$ with $\theta_i \sim \text{Unif}[0, 2\pi]$ and then evaluating the trace of the corresponding Hessian matrix $\text{Tr}[H(\theta_m)]$. These values are used to build the histogram showing the frequency distribution of $\text{Tr}[H]$. The length of the arrows are, respectively: “Numerical 2σ ” (black solid line) twice the numerical standard deviation, “Approximation” (dashed red) twice the square root of the approximation in Eq. (31), “Bound” (dashed-dotted green) twice the square root of the upper bound in Eq. (30)

The circuit used has $M = 92$ parameters and $n = 4$ qubits, and plugging these values in Equation (31) yields a standard deviation $\sigma_U = \text{Std}_U[\text{Tr}[H]] \approx 11$. Then, if the behaviour of the PQCs in practical scenarios is well described by its random parameter regime, one expects the trace of the Hessian to be on average zero and in general not much bigger (in absolute value) than $\sigma_U \approx 11$. With this order of magnitude for the trace, the first order correction Equation (28) even with a Gaussian noise level of $\sigma = 0.1$ is very small, as it amounts to

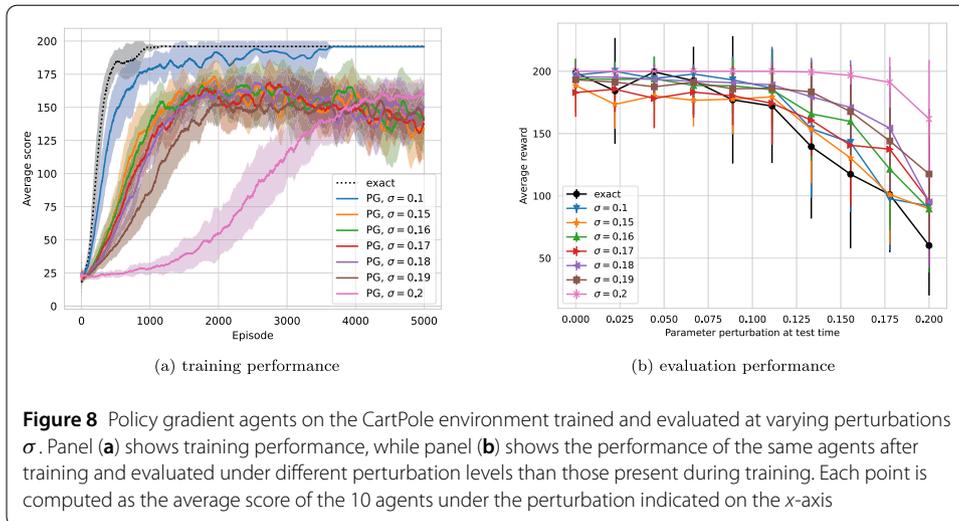
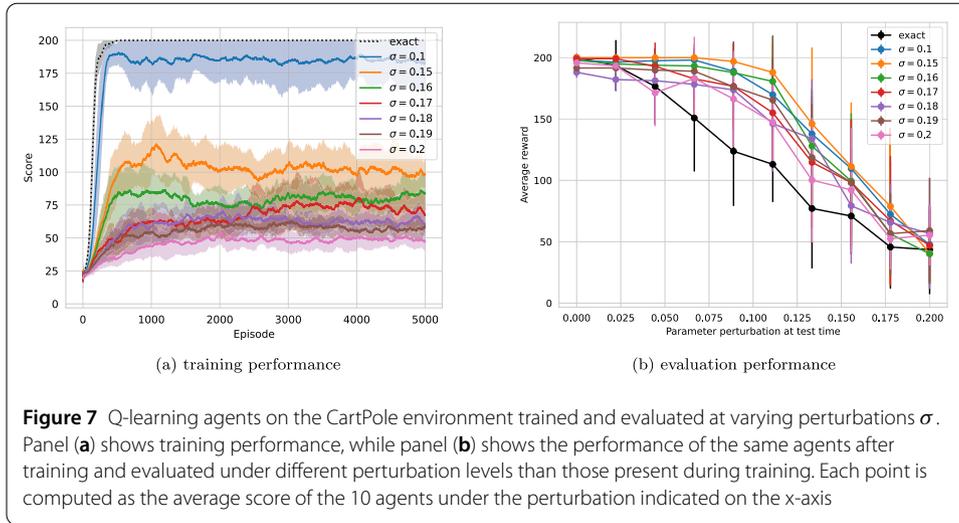
$$|\mathbb{E}[f(\theta + \delta\theta)] - f(\theta)| \approx \frac{1}{2}\sigma^2 |\text{Tr}[H(\theta)]| \approx 0.05.$$

Summing up, for those PQCs whose cost landscape is close to being flat, then Gaussian perturbations on the variational parameters will have a limited impact on the output of the quantum circuit.

5.3 Numerical results

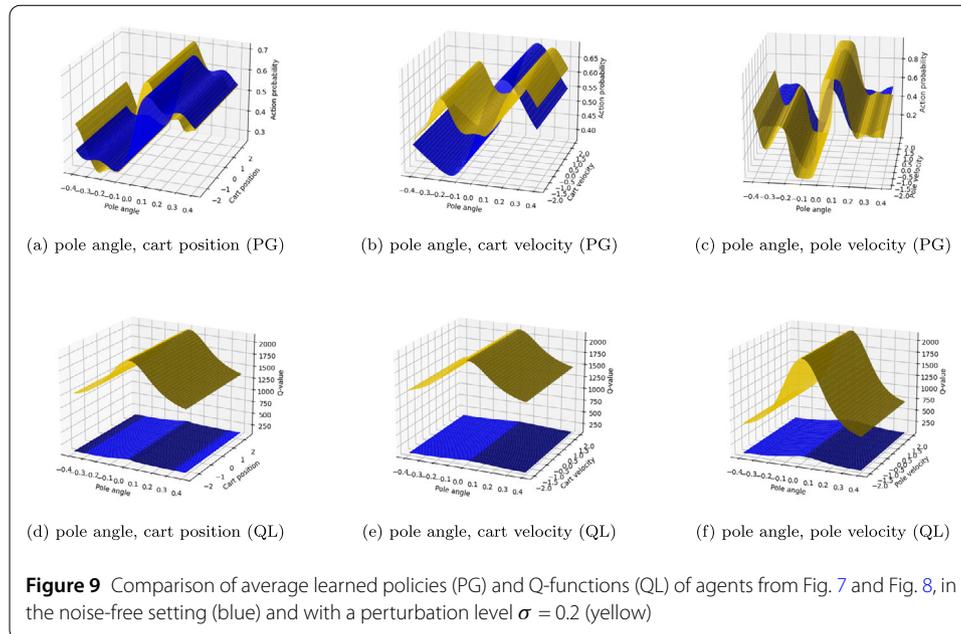
5.3.1 CartPole

First, we evaluate the performance of policy gradient and Q-learning algorithms when Gaussian perturbations are applied at each circuit evaluation during training. In Fig. 7 (a) and (b), we show the training and evaluation performance, respectively, of Q-learning agents in the CartPole environment with perturbations in the range $\sigma \in \{0, 0.1, 0.15, 0.16, 0.17, 0.18, 0.19, 0.2\}$. Only the agent trained with noise level $\sigma = 0.1$ learns the environment successfully and remains close to optimal performance. As suggested by our theoretical analysis in Sect. 5.1, performance starts to degrade as we consider higher perturbations of $\sigma > 0.1$, and none of those agents manage to achieve a better performance than a score of 125 on average. In Fig. 7 (b) we evaluate the performance of trained agents when they act in an environment with different perturbation levels than those present when they were trained. Even agents that do not perform well during training achieve close to optimal performance when evaluated in the noise-free setting. This suggests that despite their bad



training performance due to the added perturbations, these agents still learn a good Q-function. Notably, the agents trained without noise perform worst when they are evaluated under various levels of perturbations.

Results for agents trained with the policy gradient method are shown in Fig. 8 (a). While again only the agents trained with a perturbation of $\sigma = 0.1$ perform well and even reach optimal performance, agents with higher perturbations also largely stay close to optimal performance with a final score of 125 on average. Even the agent trained with a relatively high $\sigma = 0.2$ is robust in this setting, even though it requires by far the most training episodes to get to a good score. This positive trend is also visible in Fig. 8(b), where we see that all agents achieve close to optimal performance when evaluated with perturbation levels $\sigma \leq 0.1$, which is again in line with our theoretical analysis in Sect. 5.1. The difference between agents trained with Gaussian perturbations and those trained without is not as large as in the Q-learning setting, and at evaluation time both algorithms perform similarly. Another observation about the policy gradient agents is that those trained with $\sigma = 0.2$ achieve optimal or close to optimal performance in the environment under various perturbation levels at evaluation time, and are the most robust out of all agents trained



in this setting. Overall, the policy gradient method shows a larger resilience to Gaussian noise in our experiments for the CartPole environment. It is an open question why this is the case, however, we did not observe better performance of the policy gradient algorithm under noise in general, as results in later sections will show.

In addition to studying the performance of Q-learning and policy gradient agents at training and evaluation time, we visualize the learned policies and Q-functions of both in the noisy and noise-free setting in Fig. 9. As learned policies and Q-functions can look different even when training the same agent twice, we show averages of the ten agents shown in Fig. 7 and Fig. 8 for both algorithms, and for perturbation levels of $\sigma = 0$ (blue) and $\sigma = 0.2$ (yellow), respectively. The CartPole environment has four inputs: cart position and velocity, and pole angle and velocity. To visualize the learned policies and Q-functions, we show the probabilities and Q-values for taking the action “right” as a function of pairs of state values. The state inputs that are not in the figure are set to zero, and for the sake of clarity we do not apply perturbations to the parameters when visualizing the policy. In Fig. 9 (a)-(c), we see results for policy gradient agents. Overall, it can be seen that the agents trained without perturbations learn smoother policies, hence for most states there is a clear decision on which action to take. Training with perturbations makes the policies slightly more rippled, but they still mostly follow the contours of the policy learned under ideal conditions.

The approximated Q-functions can be seen in Fig. 9(d)-(f). One observation we make here is that the range that Q-values take blows up considerably compared to the noise-free setting. This is due to the trainable output weights that the expectation values are multiplied with in the Q-learning setting (see Sect. 3) becoming considerably larger for agents trained in the noisy setting. However, as we can see in the Appendix in Fig. 17, the shapes of the learned Q-functions of the noise-free and noisy agents are still very similar, which explains why even the agents trained with $\sigma = 0.2$ perform almost optimally when evaluated without perturbations in Fig. 7 (b). We also note that the range of Q-values of both the noisy and noise-free agents is much larger than the range of optimal Q-values

given in [29]. This can be understood as the agent consistently overestimating the expected return, a problem known to arise in classical Q-learning, and which is exacerbated by noise [59]. However, the authors of [29] also point out that in the function approximation setting, it is more important to learn the order of Q-values for each state (i.e., preserving that the argmax Q-value corresponds to the optimal action) than learning a close representation of the optimal Q-values.

5.3.2 TSP

In this section, we study the performance of Q-learning and policy gradient algorithms with Gaussian coherent noise in the TSP environment. Panels (a) and (b) in Fig. 10 show the training and evaluation performance of Q-learning agents in this environment under perturbations in the range $\sigma \in \{0, 0.1, 0.15, 0.16, 0.17, 0.18, 0.19, 0.2\}$. We note that the Q-learning agents trained without noise already converge after 600 episodes on average, but to get an equal runtime in terms of episodes for all settings, we also let them run for 10,000 episodes. This unnecessarily long runtime causes the optimizer to leave the local minimum again, which we ignore as an artifact here and consider the lowest average approximation ratio for the comparison with the other models.

For the TSP environment, we observe that with increasing levels of Gaussian perturbations, convergence of agents is delayed and their final approximation ratio becomes worse compared to the noise-free agents' performance. Still, all agents seem to learn very similar policies despite being trained with different settings of σ , as we can see by their almost identical performance at evaluation time shown in Fig. 10 (b). Despite a drop in performance during training, the final performance of the models on a test set of previously unseen TSP instances stays almost unaffected by the noise present during training. While we see that agents trained with more noise seem to learn more noise-robust policies as in the case of the CartPole environment, this effect is not as pronounced here. Additionally, we again see that performance of trained models in Fig. 10 (b) starts to drop at $\sigma > 0.1$, as indicated by our theoretical analysis in Sect. 5.1. While the policy gradient method shows a certain robustness to noise during training in the CartPole environment, this is not the case for the TSP environment, as we show in Fig. 11 (a). The only agent that gets close in performance to the noise free agent is the one trained with $\sigma = 0.1$, while higher perturbations yield agents that are relatively bad with an approximation ratio between 1.4 and 1.6

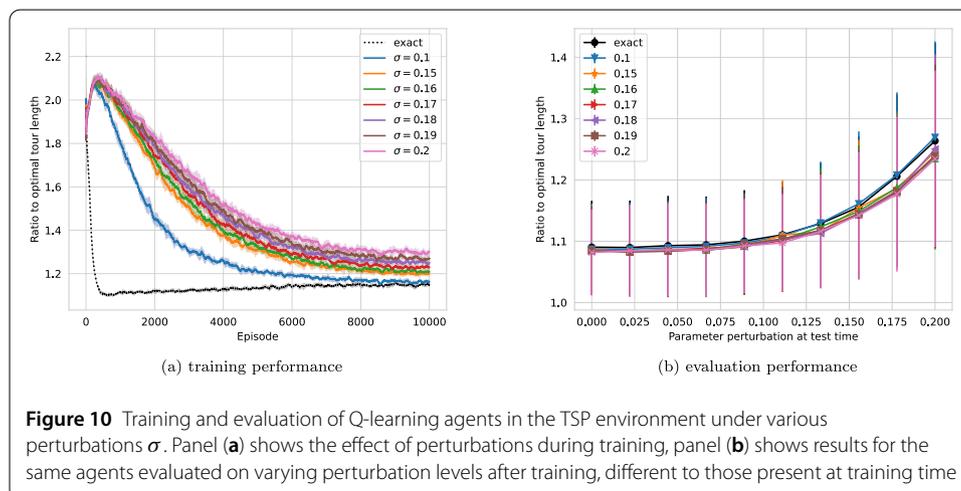


Figure 10 Training and evaluation of Q-learning agents in the TSP environment under various perturbations σ . Panel (a) shows the effect of perturbations during training, panel (b) shows results for the same agents evaluated on varying perturbation levels after training, different to those present at training time

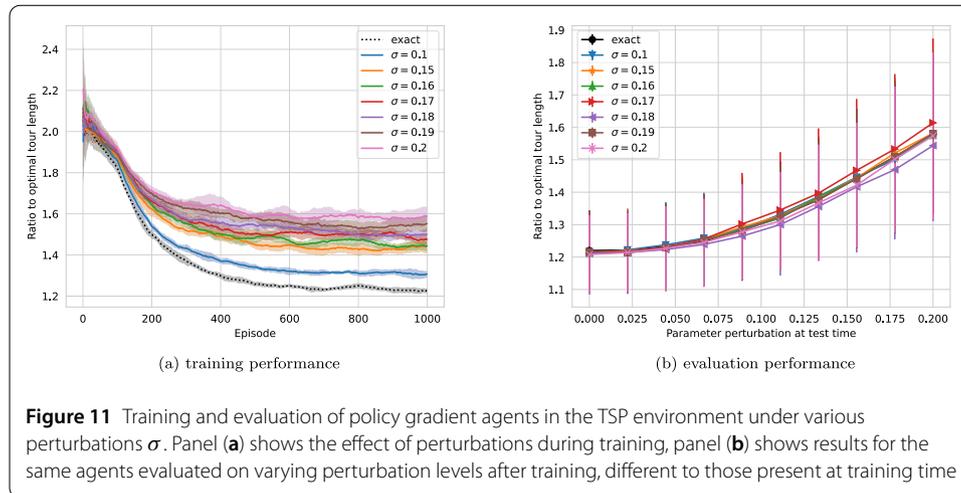


Figure 11 Training and evaluation of policy gradient agents in the TSP environment under various perturbations σ . Panel (a) shows the effect of perturbations during training, panel (b) shows results for the same agents evaluated on varying perturbation levels after training, different to those present at training time

on average. However, again, all agents seem to learn similar policies as indicated by their test performance in Fig. 11 (b). Similar to CartPole, the agents' performance on the test set under varying perturbation levels closely matches that of the noise-free agents, and again we see a large drop in performance for perturbations that are higher than $\sigma = 0.1$.

Overall, the Q-learning algorithm performs better in the TSP environment than the policy gradient method. The optimal tour for each TSP instance is deterministic, so using a stochastic policy as in the policy gradient approach introduces an additional source of error, as there is always a non-zero probability to choose a non-optimal action. This leads to an increased susceptibility to the Gaussian perturbations present during the evaluation of the policy gradient algorithm. This is not the case for Q-learning, where choices are made based on the argmax Q-value. Additionally, the ansatz that we use does not separate between data encoding and trainable parameters as described in Sect. 3. As the optimal tour of a TSP instance does not change upon small perturbations of the edge weights, this leads to a relative robustness of this ansatz used in conjunction with Q-learning to Gaussian coherent noise in this environment.

6 Incoherent noise

The Gaussian perturbation noise that we studied in Sect. 5 is well-suited to model coherent errors due to imprecision in the control of the quantum device, but it does not reflect noise that results from undesired interactions of the quantum system with its environment. To study the effect of this type of incoherent noise we perform additional experiments in this section.

We simulate this type of noise with TensorFlow Quantum (TFQ) [60], where they are implemented through a Monte-Carlo trajectory sampling method [61, 62] that approximates the effect of noise by averaging over state vectors generated from a probabilistic application of the noise channel. This method of simulating noise essentially trades off the overhead in memory needed to store the $2^n \times 2^n$ sized density matrices necessary to simulate incoherent noise, with a runtime overhead. The precision of this approximation is determined by the number of repetitions, which specifies how many "noisy" state vectors are used. This adds a stochastic element to the simulation of the noise channels, and we get closer to simulating the exact noise model as the number of trajectories increases. Depending on the environments, we choose the number of trajectories so that it is possible

to perform simulations in a reasonable time frame, and specify this number individually for each of the experiments below. We note that the runtime requirements for CartPole when simulating this type of noise are especially high, as the number of time steps in each episode, as well as the number of episodes itself depends strongly on the performance of the agent. In particular, agents that perform neither very well nor very badly, which are exactly the noise configurations we are interested in studying here, take especially long to simulate, as they do not converge early by solving the environment, but still take on the order of 100 time steps in each episode. Therefore we focus our attention mainly on the TSP environment in this section.

6.1 Depolarizing noise

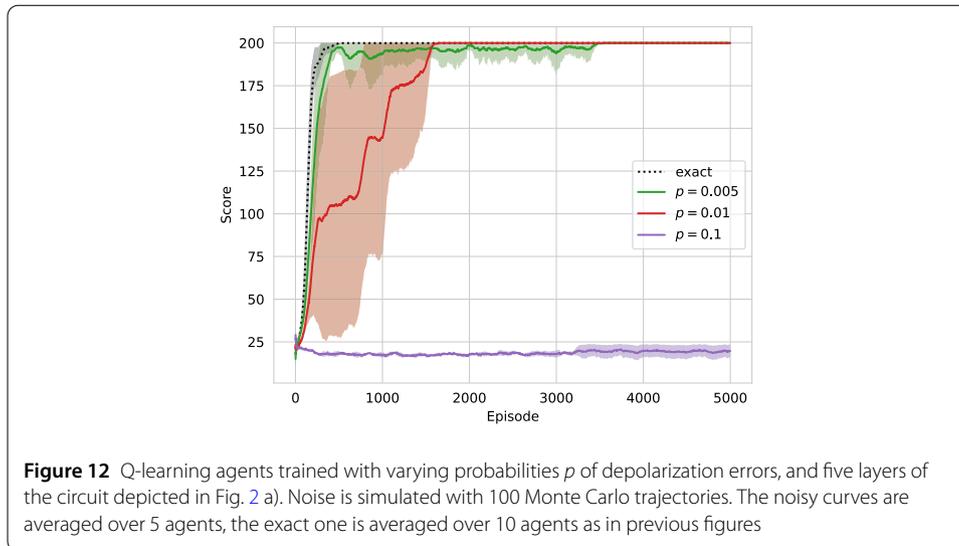
Depolarization noise affects a quantum state by either replacing it with the completely mixed state with probability p , or leaving it untouched otherwise [63]. Let ρ be the density matrix of a qubit, then depolarizing noise is defined by the map

$$\mathcal{D}_p(\rho) = (1 - p)\rho + p\frac{\mathbb{1}}{2}. \quad (32)$$

We model depolarization noise with Cirq [61] and TFQ by appending a layer of local depolarizing channels to every qubit after each time step of the computation, where a time step is defined as the largest set of gates that can be implemented simultaneously. This implementation takes into account the possibility of cross-talk between qubits [64]. Also, note that while the use of depolarizing channels alone may not be a good approximation of real single qubits errors, it may become a good effective description of the overall noise process for the case where many qubits and layers are used [65].

In our simulations, we assume that both single- and two-qubits gates are noisy, and consist of a composition of the ideal gates followed by *local* depolarizing channels of equal probability p , acting independently on each qubit. In particular, the application of a depolarizing noise channel is implemented by performing one out of four actions at each circuit execution (trajectory): do nothing with probability $1 - p$, or apply at random one of the three Pauli operators with probability p , and then average over the results. We remark that the average gate error of single-qubit gates in currently available superconducting quantum computing hardware is of the order of $r \lesssim 0.01$, with gate fidelities exceeding $> 99\%$. Finally, we note that one can relate the depolarisation strength p to the average gate error r over single qubit Cliffords, as measured by Randomized Benchmarking (RB) [64, 66, 67] and commonly reported for quantum devices [68, 69], via $r = p/2$. However, our circuits do not only use Cliffords, and moreover the RB's estimates for the gate error depend on the basis gates available on the device. Therefore, one should consider our simulations with depolarizing noise of strength p as a proxy for a quantum device whose average error rate r is of the same order of magnitude of p . While a single-qubit error noise model may not be accurate enough to closely mimic the behaviour of a real quantum device, it gives us the possibility to study the effect of single-qubit errors separately, before we go on to study a noise model that also includes two-qubit gate errors in Sect. 6.2.

As mentioned above, simulating incoherent noise has high runtime requirements, so in the following we limit our studies to: (i) Q-learning in the CartPole environment, and (ii) the policy gradient method in the TSP environment. We pick these settings as they were the ones that were more sensitive to Gaussian coherent noise in our studies in Sect. 5,



and in that sense represent the worst case instances from the previous section. To simulate the noisy quantum circuits, we use the Monte Carlo sampling as described above, where the number of trajectories used depends on the environment. As the CartPole environment requires a very high number of environment interactions (the better the agent, the more circuit evaluations are required per episode), we use 100 trajectories in this setting. In the TSP environment, the number of steps in each episode is constant and therefore we can use a higher number of 1000 trajectories and still perform simulations in a timely manner.

Figure 12 shows results of Q-learning agents trained in the CartPole environment with various error probabilities p . Agents with a realistic error probability of up to $p = 0.01$ still solve the environment in less than 2000 episodes on average. Agents trained with error probability $p = 0.005$ reach higher scores almost as quickly as agents trained in the noise-free setting, but stay somewhat unstable until they solve the environment after 3500 episodes on average. When the noise probability is increased to $p = 0.1$, we see that agents fail to make any learning progress at all.

Figure 13 shows the performance of the policy gradient method under one-qubit depolarization errors in the TSP environment. In this setting, agents trained with error probability $p = 0.01$, as is a realistic assumption on current devices, perform noticeably worse than agents in the noise-free setting with a drop in approximation ratio of around 0.2 on average. Only when we consider an error probability of $p = 0.001$ do we get performance that is almost exactly the same as that in the noise-free case. Similar to the results of the Q-learning agent in the CartPole environment, agents trained with an error probability of $p = 0.1$ show no meaningful learning progress.

6.2 Noise model based on current hardware

After studying the effect of single-qubit depolarization errors in Sect. 6.1, we now study the performance of the Q-learning algorithm in the TSP environment in the presence of a more realistic noise model that captures the behaviour of a near-term superconductive quantum device. The error sources we incorporate into this noise model are the following: single-qubit and two-qubit depolarization errors, single qubit amplitude damping error,

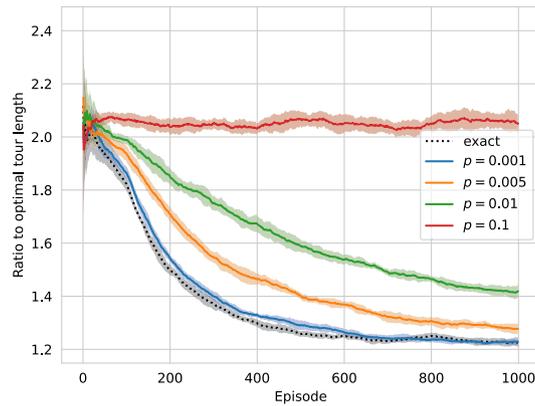


Figure 13 Policy gradient agents trained in the TSP environment with varying probabilities p of depolarization error, with one layer of the circuit depicted in Fig. 2 c). Noise is simulated with 1000 Monte Carlo trajectories. All curves are averaged over 10 agents

and measurement noise. While hardware providers like IBM and Google offer the possibility of simulating noise models of specific devices, we do not want to take device-specific factors like qubit topology and native gate sets into account in this work, as the performance in these settings also depends strongly on the quality of the circuit compiled to the native gate set and qubit connectivity [70]. Instead, we define a custom noise model based on gate fidelities published by hardware vendors, but do not take the above details into account. To determine realistic settings for the error probability of each noise source, we use calibration data published by IBM [71] at the time of writing. The noise model used in our simulation is specified as follows:

- *Depolarization error*: Single qubit depolarization channels with $p = 0.001$ are applied after every single qubit gate. Two-qubit depolarization errors, defined by properly adjusting the definition in Equation (32), with $p_2 = 0.01$ are applied after every two-qubit gate on the corresponding pair of qubits.
- *Amplitude damping error*: Amplitude damping channels with decay parameter $\gamma = 0.003$ are applied after each single- and two-qubit gate on the corresponding qubits. Such a decay rate is valid for real devices having single qubit gate durations of $t = 35$ ns, and average qubit decay times $T_1 \approx 100$ μ s, which correspond to a decay parameter of $\gamma = 1 - \exp(-t/T_1) \approx 0.0003$.
- *Measurement noise* Measurement errors are modeled by appending a bit-flip channel with probability $p = 0.01$ to every qubit right before the measurement process.

We recall that the circuit ansatz for the TSP environment is the one depicted in Fig. 2(c), where input information about the edge weights of the TSP instance is encoded by means of two-qubit gates. We therefore chose to study this ansatz in the context of a noise model that incorporates two-qubit errors, as we expect that these types of errors will affect performance of an ansatz that encodes crucial information in two-qubit gates more severely. Additionally, it is hard to perform simulations in this setting for the CartPole environment in a reasonable amount of time, as discussed above. For these reasons, we restrict our attention to the TSP environment in this section.

Figure 14 shows results averaged over five Q-learning agents in the TSP environment for each of the error probability configurations of the custom noise model described above.

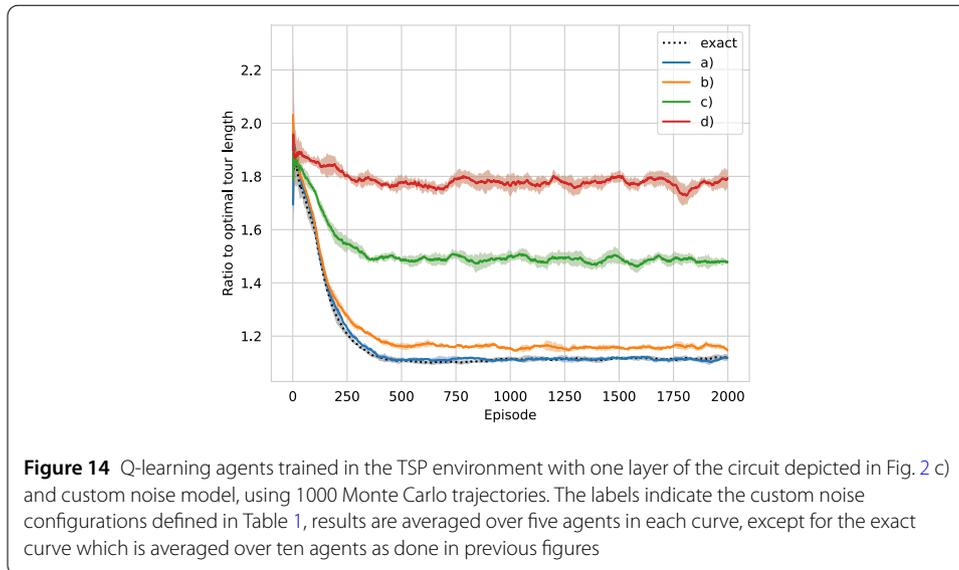


Table 1 Error strengths for the configurations of the custom noise model used in Fig. 14.

Depolarization (1Q) indicates the single qubit depolarising channel applied after each single-qubit gate, and similarly for 2Q for two-qubit gates. Configuration a) in bold is based on error rates published by IBM at the time of writing, as described in the main text

Error source	a)	b)	c)	d)
Depolarization (1Q)	0.001	0.001	0.01	0.1
Depolarization (2Q)	0.01	0.01	0.1	0.2
Amplitude damping	0.0003	0.03	0.03	0.1
Bitflip (measurement)	0.01	0.01	0.1	0.1

We show the specific error probabilities used for the simulations in Table 1. Configuration a) corresponds to error probabilities that are consistent with those present on current quantum hardware as described above. Based on this, we specify three other error probabilities b) - d) by increasing the error on varying error sources. We note that while the error probabilities themselves in configuration a) are consistent with those on current hardware, our simulation is only an approximation of this error due to the Monte Carlo trajectory sampling method described in Sect. 6. To perform simulations in a reasonable time frame, we use 1000 trajectories for each circuit evaluation. The circuit that we simulate has 145 gates (counting a ZZ-gate as two CNOTs and one Z gate), and for small error probabilities the chance of applying each of the noise channels is relatively small. This means that in each trajectory, a relatively small number of noise channels is applied. Hence we expect that the results in Fig. 14 are slightly better than what we would get if the exact noise model was simulated (i.e., in the limit of a large number of trajectories, or by considering the full density matrix).

Looking at the results in Fig. 14, we see that for configuration a) (blue), the performance of the agents matches those of the noise-free ones (dotted black) almost exactly, and the noise model based on realistic error strengths of current devices does not affect training. We see a slight drop in performance when we increase the error probability of the amplitude damping channels from 0.0003 to 0.03 (orange), as described in Table 1, column b). For configuration c), we also increase the other remaining error sources' probabilities, which leads to a considerable drop in performance. In configuration d), we assume ex-

tremely high error probabilities for each of the noise channels, which leads to a complete failure of the agents to make any meaningful learning progress in this environment.

7 Conclusions

Our goal in this work was to evaluate the resilience of variational RL algorithms to various types of noise that are present on real quantum hardware. First, we investigated shot noise, which results from the probabilistic nature of quantum measurements. We introduced a method to reduce the number of shots to train a Q-learning agent, motivated by the specific structure of the underlying RL algorithm. Our shot allocation technique enables a more shot-frugal training of variational Q-learning models with little or no effect on the final performance of the agents.

After considering shot noise, we moved on to study the effect of Gaussian coherent errors that can arise on real hardware due to miscalibration of the device, or imprecise pulse sequences that implement the parameterised gates in the quantum circuit. We gave an analytic expression for how this type of noise affects the output of a quantum RL agent, and provided a bound on the standard deviation of the Gaussian error that elucidates the tolerable magnitude of the error on the output of a quantum model. We confirm this bound in our simulations, where we study the effect of various levels of Gaussian perturbations on the performance of training policy gradient and Q-learning agents in two different environments. For one of these environments, we find that agents trained with higher noise probabilities also learn more robust policies and Q-functions, in the sense that under evaluation of different perturbation levels, these agents achieve optimal or close to optimal performance more often.

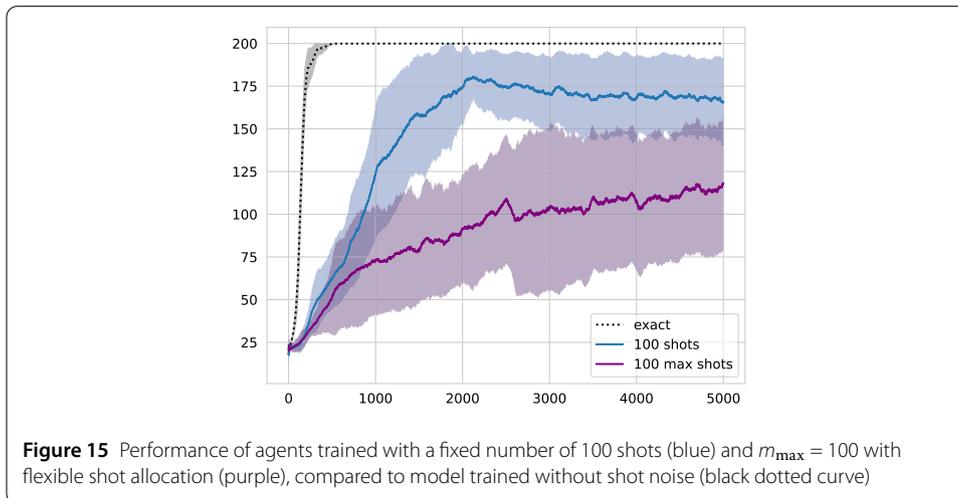
Finally, we studied incoherent noise that emerges in real hardware due to undesired interactions of the qubits with the surrounding environment, as the device is not completely shielded from external effects. To this end, we consider single-qubit depolarization errors, as well as a custom noise model that combines single- and two qubit depolarization errors, amplitude damping errors, and bitflip (measurement) errors. For the latter, we perform simulations with realistic error probabilities for each of the noise channels, in line with data published for IBM devices at the time of writing.

Overall, we find that the effect of noise on training variational RL algorithms for Q-learning and the policy gradient method depends strongly on the strength of the noise, as well as the type of noise itself. For some cases, like decoherence errors with realistic error probabilities of current devices, the drop in performance is relatively small. On the other hand, we find that large Gaussian perturbations as well as errors induced by the probabilistic nature of quantum measurements can affect performance in highly detrimental ways. Additionally, we find that for Gaussian coherent noise agents that are trained with higher perturbations learn more noise-robust policies in some cases, similar to results in classical literature, where noise is used as a regularization technique.

While our results were performed in a regime that is still efficiently simulable on classical computers, it is an interesting question for future work to consider the implications of noise-robustness of large-scale quantum models in light of recent results which show that in certain settings, the outputs of noisy quantum circuits can be efficiently approximated classically [72, 73]. This raises the question to what extent an inherent noise-robustness of hybrid variational quantum machine learning affects the possibility to achieve a quantum advantage with these types of models.

On the practical side, the optimization procedures that we used in this work were the same as those commonly used to train models in noise-free simulations and are not tailored to account for quantum hardware specific noise. This raises the question on how optimization methods that are tailored for the special characteristics of variational quantum models could further improve the performance of these types of models in a noisy setting. For the optimization of PQC parameters in the combinatorial optimization or quantum chemistry setting, it is known that some optimization methods, like simultaneous perturbation stochastic approximation (SPSA), actually become better with noise. It is an interesting area of future research to design quantum-specific optimization routines for machine learning that address or even combat specific types of noise, for example leveraging effective quantum error mitigation techniques [74–76]. Our work motivates the study of these types of optimization methods, as well as continued efforts to find learning tasks where variational RL algorithms can potentially provide an advantage.

Appendix A: Additional results for flexible vs. fixed number of shots in Q-learning



Appendix B: Gaussian noise analysis

In this Appendix we perform the noise analysis of a scalar function whose parameters are corrupted by independently distributed Gaussian perturbations. Let $f : \mathbb{R}^M \rightarrow \mathbb{R}$ be the function under investigation, whose parameters $\theta = (\theta_1, \dots, \theta_M) \in \mathbb{R}^M$ are corrupted by a Gaussian noise $\theta_i \rightarrow \theta_i + \delta\theta_i$ with zero mean and variance σ^2 , i.e.

$$\begin{aligned} \delta\theta_i &\sim \mathcal{N}(0, \sigma^2) \quad \forall i = 1, \dots, M, \\ \mathbb{E}[\delta\theta_i] &= 0, \\ \mathbb{E}[\delta\theta_i \delta\theta_j] &= \sigma^2 \delta_{ij}. \end{aligned} \tag{B.1}$$

Since the perturbations are independently distributed and Gaussian, all higher order moments can be evaluated starting from two points correlators of the form $\mathbb{E}[\delta\theta_i \delta\theta_j]$, as

dictated by *Wick’s formulas* for multivariate normal distributions

$$\begin{aligned} \mathbb{E}[\delta\theta_{i_1} \cdots \delta\theta_{i_{2n+1}}] &= 0, \\ \mathbb{E}[\delta\theta_{i_1} \cdots \delta\theta_{i_{2n}}] &= \sum_{\mathcal{P}} \mathbb{E}[\delta\theta_{k_1} \delta\theta_{k_2}] \cdots \mathbb{E}[\delta\theta_{k_{2n-1}} \delta\theta_{k_{2n}}], \end{aligned} \tag{B.2}$$

where with \mathcal{P} we denote all the possible distinct $(2n - 1)!!$ pairings of the n variables, as these can be used to express all higher order even moments in terms of products of second moments. Note that all the terms involving an odd number of perturbations $\delta\theta_i$ vanish, and only even moments of remain. For example, expression (B.2) for the fourth-order moment ($n = 4$) amounts to

$$\begin{aligned} \mathbb{E}[\delta\theta_i \delta\theta_j \delta\theta_k \delta\theta_m] &= \mathbb{E}[\delta\theta_i \delta\theta_j] \mathbb{E}[\delta\theta_k \delta\theta_m] + \mathbb{E}[\delta\theta_i \delta\theta_k] \mathbb{E}[\delta\theta_j \delta\theta_m] + \mathbb{E}[\delta\theta_i \delta\theta_m] \mathbb{E}[\delta\theta_j \delta\theta_k] \\ &= \sigma^4 (\delta_{ij} \delta_{km} + \delta_{ik} \delta_{jm} + \delta_{im} \delta_{jk}). \end{aligned} \tag{B.3}$$

We now proceed considering the multi dimensional Taylor expansion of the function $f(\boldsymbol{\theta} + \delta\boldsymbol{\theta})$ around the noise-free point. Up to arbitrary order, this reads

$$\begin{aligned} f(\boldsymbol{\theta} + \delta\boldsymbol{\theta}) &= f(\boldsymbol{\theta}) + \sum_{i=1}^M \frac{\partial f(\boldsymbol{\theta})}{\partial \theta_i} \delta\theta_i + \frac{1}{2!} \sum_{i,j=1}^M \frac{\partial^2 f(\boldsymbol{\theta})}{\partial \theta_i \partial \theta_j} \delta\theta_i \delta\theta_j \\ &+ \frac{1}{3!} \sum_{i,j,k=1}^M \frac{\partial^3 f(\boldsymbol{\theta})}{\partial \theta_i \partial \theta_j \partial \theta_k} \delta\theta_i \delta\theta_j \delta\theta_k + \cdots, \end{aligned} \tag{B.4}$$

where we used the equal sign because we are considering the full Taylor series, and we assume that this converges to the true function (this statement can be made precise by showing that the reminder term of the expansion goes to zero as the order of expansion goes to infinity).

Before proceeding, we simplify the notation to make the calculation of the Taylor expansion easier to follow. First, we denote the partial derivatives with respect to parameter θ_i as $\partial_i := \partial/\partial\theta_i$, and similarly for higher order derivatives, for example $\partial_{ij} = \partial^2/\partial\theta_i \partial\theta_j$. Also, we suppress the explicit dependence of the function on $\boldsymbol{\theta}$, using the short-hand f instead of $f(\boldsymbol{\theta})$. At last, we make use of Einstein’ summation notation where repeated indexes imply summation.

With this setup, using Eqs. (B.1), (B.2) and (B.3) in (B.4), one can evaluate the expectation value of the function over the perturbations’ distributions as

$$\begin{aligned} \mathbb{E}[f(\boldsymbol{\theta} + \delta\boldsymbol{\theta})] &= f(\boldsymbol{\theta}) + \partial_i f \mathbb{E}[\delta\theta_i] + \frac{1}{2} \partial_{ij} f \mathbb{E}[\delta\theta_i \delta\theta_j] + \frac{1}{3!} \partial_{ijk} f \mathbb{E}[\delta\theta_i \delta\theta_j \delta\theta_k] \\ &+ \frac{1}{4!} \partial_{ijkl} f \mathbb{E}[\delta\theta_i \delta\theta_j \delta\theta_k \delta\theta_l] + \cdots \\ &= f(\boldsymbol{\theta}) + \frac{\sigma^2}{2} \partial_{ij} f \delta_{ij} + \frac{\sigma^4}{4!} \partial_{ijkl} f (\delta_{ij} \delta_{kl} + \delta_{ik} \delta_{jl} + \delta_{il} \delta_{jk}) + \cdots \\ &= f(\boldsymbol{\theta}) + \frac{\sigma^2}{2} \sum_i \frac{\partial^2 f}{\partial \theta_i^2} + \frac{\sigma^4}{4!} 3 \sum_{ij} \frac{\partial^4 f}{\partial \theta_i^2 \partial \theta_j^2} + \cdots, \end{aligned} \tag{B.5}$$

where in the last line we simplified the fourth order term as

$$\begin{aligned} \mathbb{E}[f^{(4)}] &= \frac{\sigma^4}{4!} \partial_{ijkm} f (\delta_{ij} \delta_{km} + \delta_{ik} \delta_{jm} + \delta_{im} \delta_{jk}) \\ &= \frac{\sigma^4}{4!} \left(\sum_{ik} \frac{\partial^4 f}{\partial \theta_i^2 \partial \theta_k^2} + \sum_{ij} \frac{\partial^4 f}{\partial \theta_i^2 \partial \theta_j^2} + \sum_{im} \frac{\partial^4 f}{\partial \theta_i^2 \partial \theta_m^2} \right) \\ &= \frac{\sigma^4}{4!} 3 \sum_{ij} \frac{\partial^4 f}{\partial \theta_i^2 \partial \theta_j^2}. \end{aligned}$$

Since the expectation values involving an odd number of perturbations vanish, only the even order terms survive, and these can be expressed as

$$\mathbb{E}[f^{(2n)}] = \frac{\sigma^{2n}}{(2n)!} (2n-1)!! \sum_{i_1, \dots, i_n} \frac{\partial^{2n} f(\boldsymbol{\theta})}{\partial \theta_{i_1}^2 \dots \partial \theta_{i_n}^2}, \tag{B.6}$$

where the coefficient $(2n-1)!!$ is the number of distinct pairings of $2n$ objects, which comes from Eq. (B.1).

Thus, the full Taylor series can be formally written as

$$\mathbb{E}[f(\boldsymbol{\theta} + \delta\boldsymbol{\theta})] = f(\boldsymbol{\theta}) + \sum_{n=1}^{\infty} \frac{\sigma^{2n}}{(2n)!} (2n-1)!! \sum_{i_1, \dots, i_n=1}^M \frac{\partial^{2n} f(\boldsymbol{\theta})}{\partial \theta_{i_1}^2 \dots \partial \theta_{i_n}^2} \tag{B.7}$$

$$= f(\boldsymbol{\theta}) + \frac{\sigma^2}{2} \text{Tr}[H(\boldsymbol{\theta})] + \sum_{n=2}^{\infty} \frac{\sigma^{2n}}{(2n)!} (2n-1)!! \sum_{i_1, \dots, i_n=1}^M \frac{\partial^{2n} f(\boldsymbol{\theta})}{\partial \theta_{i_1}^2 \dots \partial \theta_{i_n}^2}, \tag{B.8}$$

where we introduced the Hessian matrix $H(\boldsymbol{\theta})$, whose elements are given by $[H(\boldsymbol{\theta})]_{ij} = \partial_{ij} f(\boldsymbol{\theta})$, and we see that this term represent the first non-vanishing correction to the function caused by the perturbation.

Our goal is to bound the absolute error

$$\varepsilon_{\boldsymbol{\theta}} := \left| \mathbb{E}[f(\boldsymbol{\theta} + \delta\boldsymbol{\theta})] - f(\boldsymbol{\theta}) \right| = \left| \sum_{n=1}^{\infty} \frac{\sigma^{2n}}{(2n)!} (2n-1)!! \sum_{i_1, \dots, i_n=1}^M \frac{\partial^{2n} f(\boldsymbol{\theta})}{\partial \theta_{i_1}^2 \dots \partial \theta_{i_n}^2} \right| \tag{B.9}$$

caused by the Gaussian noise, and we can do that by using the property that all the derivatives of most PQC (Parametrized Quantum Circuit) are bounded. In fact, for those circuits for which a *parameter-shift* rule holds [52, 53], one can show that any derivative of the function $f(\boldsymbol{\theta}) = \langle O \rangle = \text{Tr}[OU(\boldsymbol{\theta})|0\rangle\langle 0|U^\dagger(\boldsymbol{\theta})]$ obeys

$$\left| \frac{\partial^{\alpha_1 + \dots + \alpha_M} f(\boldsymbol{\theta})}{\partial \theta_1^{\alpha_1} \dots \partial \theta_M^{\alpha_M}} \right| \leq \|O\|_{\infty}, \tag{B.10}$$

where $\|O\|_{\infty}$ is the infinity norm of the observable, namely its largest absolute eigenvalue. We give a proof of this below in Sect. B.1.

Plugging this in Eq. (B.9), we can obtain an upper bound to the error $\varepsilon_{\boldsymbol{\theta}}$ as desired. Indeed, remembering that for even numbers the double factorial can be expressed as

$(2n - 1)!! = (2n)!/(2^n n!)$, it holds

$$\begin{aligned}
 \varepsilon_\theta &= \left| \sum_{n=1}^\infty \frac{\sigma^{2n}}{(2n)!} (2n - 1)!! \sum_{i_1, \dots, i_n=1}^M \frac{\partial^{2n} f(\theta)}{\partial \theta_{i_1}^2 \dots \partial \theta_{i_n}^2} \right| \\
 &\leq \sum_{n=1}^\infty \frac{\sigma^{2n}}{(2n)!} (2n - 1)!! \sum_{i_1, \dots, i_n=1}^M \underbrace{\left| \frac{\partial^{2n} f(\theta)}{\partial \theta_{i_1}^2 \dots \partial \theta_{i_n}^2} \right|}_{\leq \|O\|_\infty} \\
 &\leq \sum_{n=1}^\infty \frac{\sigma^{2n}}{(2n)!} (2n - 1)!! \|O\|_\infty M^n = \|O\|_\infty \sum_{n=1}^\infty \frac{1}{(2n)!} \frac{(2n)!}{2^n n!} (\sigma^2 M)^n \\
 &= \|O\|_\infty \sum_{n=1}^\infty \frac{(M\sigma^2/2)^n}{n!} = \|O\|_\infty (e^{\sigma^2 M/2} - 1) \\
 \implies \varepsilon_\theta &= |\mathbb{E}[f(\theta + \delta\theta)] - f(\theta)| \leq \|O\|_\infty (e^{M\sigma^2/2} - 1), \tag{B.11}
 \end{aligned}$$

where in the last line we used the definition of the exponential function $e^x = \sum_{n=0}^\infty \frac{x^n}{n!}$.

One can see that the noise variance σ^2 must scale as the inverse of the number of parameters $\sigma^2 \in \mathcal{O}(M^{-1})$ in order to have small deviations induced by the noise. Also, note that since the difference between the noise-free function $f(\theta)$ and its perturbed version $f(\theta + \delta\theta)$ cannot be larger than twice the maximum eigenvalue of O , $|f(\theta + \delta\theta) - f(\theta)| \leq |f(\theta + \delta\theta)| + |f(\theta)| = 2\|O\|_\infty$, the bound (B.11) is informative only as long as $\exp[M\sigma^2/2] - 1 < 2$.

It is worth noticing that an identical procedure can be used to bound the average error obtained by approximating the perturbed function with its first non-vanishing correction given by the Hessian. Indeed, starting from Eq. (B.8) are repeating the same calculation from above, one obtains

$$\left| \mathbb{E}[f(\theta + \delta\theta)] - f(\theta) - \frac{\sigma^2}{2} \text{Tr}[H(\theta)] \right| \leq \|O\|_\infty \left(e^{M\sigma^2/2} - 1 - \frac{M\sigma^2}{2} \right). \tag{B.12}$$

B.1 Parameter-shift rule and bounds to the derivatives

Let $f(\theta) = \text{Tr}[OU(\theta)|0\rangle\langle 0|U^\dagger(\theta)]$ be the expectation value of an observable O on the parametrized state $|\psi(\theta)\rangle = U(\theta)|0\rangle$ obtained with a parametrized quantum circuit $U(\theta)$. When the variational parameters $\theta \in \mathbb{R}^M$ enter in the quantum circuit via rotation gates of the form $V(\theta_i) = \exp[-i\theta_i P/2]$ with $P^2 = \mathbb{1}$ being Pauli operators, then the *parameter-shift* rule can be used to evaluate gradients of the expectation value as [52, 53]

$$\frac{\partial f(\theta)}{\partial \theta_i} = \frac{1}{2} \left(f\left(\theta + \frac{\pi}{2} \mathbf{e}_i\right) - f\left(\theta - \frac{\pi}{2} \mathbf{e}_i\right) \right), \tag{B.13}$$

where \mathbf{e}_i is the unit vector with zero entries and a one in the i -th position corresponding to angle θ_i . Similarly, by applying the parameter-shift rule twice one can express second order derivatives as follows using four evaluations of the circuit [35, 77]

$$\frac{\partial^2 f(\theta)}{\partial \theta_i \partial \theta_j} = \frac{1}{2} \left[\frac{\partial}{\partial \theta_i} f\left(\theta + \frac{\pi}{2} \mathbf{e}_j\right) - \frac{\partial}{\partial \theta_i} f\left(\theta - \frac{\pi}{2} \mathbf{e}_j\right) \right] \tag{B.14}$$

$$\begin{aligned}
 &= \frac{1}{4} \left[f\left(\boldsymbol{\theta} + \frac{\pi}{2} \mathbf{e}_j + \frac{\pi}{2} \mathbf{e}_i\right) - f\left(\boldsymbol{\theta} + \frac{\pi}{2} \mathbf{e}_j - \frac{\pi}{2} \mathbf{e}_i\right) \right. \\
 &\quad \left. - f\left(\boldsymbol{\theta} - \frac{\pi}{2} \mathbf{e}_j + \frac{\pi}{2} \mathbf{e}_i\right) + f\left(\boldsymbol{\theta} - \frac{\pi}{2} \mathbf{e}_j - \frac{\pi}{2} \mathbf{e}_i\right) \right]. \tag{B.15}
 \end{aligned}$$

In particular, for the diagonal elements $i = j$, one has

$$\begin{aligned}
 \frac{\partial^2 f(\boldsymbol{\theta})}{\partial \theta_i^2} &= \frac{1}{4} [f(\boldsymbol{\theta} + \pi \mathbf{e}_i) - 2f(\boldsymbol{\theta}) + f(\boldsymbol{\theta} - \pi \mathbf{e}_i)] \\
 &= \frac{1}{2} [f(\boldsymbol{\theta} + \pi \mathbf{e}_i) - f(\boldsymbol{\theta})], \tag{B.16}
 \end{aligned}$$

where we used the fact that $f(\boldsymbol{\theta} + \pi \mathbf{e}_i) = f(\boldsymbol{\theta} - \pi \mathbf{e}_i)$. This last equality can be seen intuitively from the 2π periodicity of the rotation gates or by direct evaluation. In fact, let $U(\boldsymbol{\theta}) = U_2 \exp[-i\theta_i P_i/2] U_1$ be a factorization of the parametrized unitary where we isolated the dependence on the parameter θ_i to be shifted. Then, since $\exp[-i2\pi P/2] = \cos \pi \mathbb{I} - i \sin \pi P = -\mathbb{I}$, one has

$$\begin{aligned}
 |\psi(\boldsymbol{\theta} - \pi \mathbf{e}_i)\rangle &= U_2 \exp[-i(\theta_i - \pi)P_i/2] U_1 |0\rangle \\
 &= U_2 \exp[-i(\theta_i - \pi)P_i/2] \underbrace{-\exp[-i2\pi P_i/2]}_{\mathbb{I}} U_1 |0\rangle \\
 &= -U_2 \exp[-i(\theta_i - \pi + 2\pi)P_i/2] U_1 |0\rangle \\
 &= -|\psi(\boldsymbol{\theta} + \pi \mathbf{e}_i)\rangle, \tag{B.17}
 \end{aligned}$$

and thus $\langle \psi(\boldsymbol{\theta} - \pi \mathbf{e}_i) | O | \psi(\boldsymbol{\theta} - \pi \mathbf{e}_i) \rangle = \langle \psi(\boldsymbol{\theta} + \pi \mathbf{e}_i) | O | \psi(\boldsymbol{\theta} + \pi \mathbf{e}_i) \rangle$.

Hence, using Eq. (B.16) it is possible to estimate the diagonal elements of the Hessian matrix with just two different evaluations of the quantum circuit.

By repeated application of the parameter-shift rule one can also evaluate arbitrary higher-order derivatives as linear combination of circuit evaluations [35, 36]. Let $\boldsymbol{\alpha} = (\alpha_1, \dots, \alpha_M) \in \mathbb{N}^M$ be a multi-index keeping track of the orders of derivatives, and let $|\boldsymbol{\alpha}| = \sum_{i=1}^M \alpha_i$. Then

$$\partial^{\boldsymbol{\alpha}} f(\boldsymbol{\theta}) := \frac{\partial^{|\boldsymbol{\alpha}|} f(\boldsymbol{\theta})}{\partial \theta_1^{\alpha_1} \dots \partial \theta_M^{\alpha_M}} = \frac{1}{2^{|\boldsymbol{\alpha}|}} \sum_{m=1}^{2^{|\boldsymbol{\alpha}|}} s_m f(\tilde{\boldsymbol{\theta}}_m), \tag{B.18}$$

where $s_m \in \{\pm 1\}$ are signs, and $\tilde{\boldsymbol{\theta}}_m$ are angles obtained by accumulation of shifts along multiple directions.

Since the output of any circuit evaluation is bounded by the infinity norm (i.e the largest absolute eigenvalue) of the observable $\|O\|_\infty = \max\{|o_i|, O = \sum_i o_i |o_i\rangle \langle o_i|\}$

$$|f(\boldsymbol{\theta})| = |\text{Tr}[O\rho(\boldsymbol{\theta})]| \leq \|O\|_\infty \|\rho(\boldsymbol{\theta})\|_1 = \|O\|_\infty \quad \forall \boldsymbol{\theta} \in \mathbb{R}^M, \tag{B.19}$$

then one can bound the sum in Eq. (B.18) simply as

$$|\partial^{\boldsymbol{\alpha}} f(\boldsymbol{\theta})| \leq \frac{1}{2^{|\boldsymbol{\alpha}|}} \sum_{m=1}^{2^{|\boldsymbol{\alpha}|}} |f(\tilde{\boldsymbol{\theta}}_m)| \leq \|O\|_\infty. \tag{B.20}$$

B.2 Average value of the Hessian of random PQC

In this section we derive the formulas (29) and (30) for the expected value of the Hessian as shown in the main text. Consider a system of n qubits and a parametrized quantum circuit with unitary $U(\theta) \in \mathcal{U}(2^n)$, where $\mathcal{U}(2^n)$ is the group of unitary matrices of dimension 2^n . Given a set of parameter vectors $\{\theta_1, \theta_2, \dots, \theta_K\}$, one can construct the corresponding set of unitaries $\mathbb{U} = \{U_1, U_2, \dots, U_K\}$, with $U_i = U(\theta_i)$ and clearly $\mathbb{U} \in \mathcal{U}(2^n)$.

It is now well known that sampling a parametrized quantum circuit from a random assignment of the parameters is approximately equal to drawing a random unitary from the Haar distribution, a phenomenon which is at the root of the insurgence of Barren Plateaus (BPs) [54–56]. Specifically, it is numerically observed that parametrized quantum circuits behave like unitary 2-designs, that is averaging over unitaries U_i sampled from \mathbb{U} yields the same result of averaging over Haar-random unitaries, up until second order moments.

As standard in the literature regarding BPs, in the following we assume that the considered parametrized unitaries (and parts of them) are indeed 2-designs, and so we make use of the following relations for integration over random unitaries [55–58, 78]

$$\mathbb{E}_U[UAU^\dagger] = \int d\mu(U)UAU^\dagger = \frac{\mathbb{1} \text{Tr}[A]}{2^n}, \tag{B.21}$$

$$\begin{aligned} \mathbb{E}_U[AUBU^\dagger CUDU^\dagger] \\ = \frac{\text{Tr}[BD] \text{Tr}[C]A + \text{Tr}[B] \text{Tr}[D]AC}{2^{2n} - 1} - \frac{\text{Tr}[BD]AC + \text{Tr}[B] \text{Tr}[C] \text{Tr}[D]A}{2^n(2^{2n} - 1)}. \end{aligned} \tag{B.22}$$

B.2.1 Statistics of the Hessian

Let $f(\theta) = \text{Tr}[OU(\theta)|0\rangle\langle 0|U(\theta)^\dagger]$ and assume that the observable O is such that $\text{Tr}[O] = 0$ and $\text{Tr}[O^2] = 2^n$, as is the case of measuring a Pauli string. As shown in Eq. (B.16), diagonal elements of the Hessian matrix H can be calculated as

$$H_{ii} = \frac{\partial^2 f(\theta)}{\partial \theta_i^2} = \frac{1}{2}[f(\theta + \pi \mathbf{e}_i) - f(\theta)]. \tag{B.23}$$

For simplicity, from now on drop the explicit dependence on the parameter vector θ when not explicitly needed. The variational parameters enter the quantum circuit via Pauli rotations $e^{-i\theta_i P_i/2}$ with $P_i = P_i^\dagger$ and $P_i^2 = \mathbb{1}$, and so the shifted unitary $U(\theta + \pi \mathbf{e}_i)$ can be rewritten as

$$U(\theta + \pi \mathbf{e}_i) = U_L e^{-i\pi P_i/2} U_R = -iU_L P_i U_R, \tag{B.24}$$

where U_L and U_R form a bipartition of the circuit at the position of the shifted angle, so that $U(\theta) = U_L U_R$.

Assuming that the set of unitaries \mathbb{U}_L generated by U_L is at least a 1-design, one has that

$$\mathbb{E}_{U_L}[f(\theta + \pi \mathbf{e}_i)] = \mathbb{E}_{U_L}[\text{Tr}[OU_L P_i U_R |0\rangle\langle 0|U_R^\dagger P_i U_L^\dagger]] \tag{B.25}$$

$$= \text{Tr}[O \mathbb{E}_{U_L}[U_L P_i U_R |0\rangle\langle 0|U_R^\dagger P_i U_L^\dagger]] \tag{B.26}$$

$$= \text{Tr}\left[O \frac{\text{Tr}[P_i U_R |0\rangle\langle 0|U_R^\dagger P_i] \mathbb{1}}{2^n}\right] = \frac{\text{Tr}[O]}{2^n} = 0, \tag{B.27}$$

where in the first line we exchanged the trace and the expectation value since both are linear operations, and in the second line we made use of Eq. (B.21) for the first moment of the Haar distribution. Similarly, one can show that if \mathbb{U}_R forms a 1-design, then averaging over it yields the same result, namely $\mathbb{E}_{\mathbb{U}_R}[f(\theta + \pi \mathbf{e}_i)] = 0$. The same calculation for $f(\theta)$ shows that $\mathbb{E}_{\mathbb{U}_R}[f(\theta)] = \mathbb{E}_{\mathbb{U}_L}[f(\theta)] = 0$.

Thus, for every diagonal element of the Hessian, if either \mathbb{U}_L or \mathbb{U}_R is a 1-design (that is Eq. (B.21) hold), then its expectation value vanishes

$$\mathbb{E}_{\mathbb{U}_R, \mathbb{U}_L}[H_{ii}] = 0 \quad \forall i \text{ if either } \mathbb{U}_L \text{ or } \mathbb{U}_R \text{ is a 1-design.} \tag{B.28}$$

The variance of the diagonal elements can be calculated in a similar manner, even though the calculation is more involved. Substituting Eq. (B.23) in the definition of the variance, one obtains

$$\begin{aligned} \text{Var}[H_{ii}] &:= \mathbb{E}[H_{ii}^2] - \mathbb{E}[H_{ii}]^2 = \mathbb{E}[H_{ii}^2] \\ &= \frac{1}{4} [\mathbb{E}[f(\theta + \pi \mathbf{e}_i)^2] + \mathbb{E}[f(\theta)^2] - 2\mathbb{E}[f(\theta + \pi \mathbf{e}_i)f(\theta)]]. \end{aligned} \tag{B.29}$$

In order to use Eq. (B.22) for second moment integrals, we can rewrite these expectation values as follow

$$\begin{aligned} \mathbb{E}[f(\theta + \pi \mathbf{e}_i)^2] &= \mathbb{E}[\text{Tr}[OU_L P_i U_R |0\rangle\langle 0| U_R^\dagger P_i U_L^\dagger]^2] \\ &= \mathbb{E}[\text{Tr}[OU_L P_i U_R |0\rangle\langle 0| U_R^\dagger P_i U_L^\dagger] \langle 0| U_R^\dagger P_i U_L^\dagger OU_L P_i U_R |0\rangle] \\ &= \mathbb{E}[\text{Tr}[OU_L P_i U_R |0\rangle\langle 0| U_R^\dagger P_i U_L^\dagger OU_L P_i U_R |0\rangle \langle 0| U_R^\dagger P_i U_L^\dagger]] \\ &= \text{Tr}[\mathbb{E}[OU_L P_i U_R |0\rangle\langle 0| U_R^\dagger P_i U_L^\dagger OU_L P_i U_R |0\rangle \langle 0| U_R^\dagger P_i U_L^\dagger]], \end{aligned} \tag{B.30}$$

and similarly for the remaining two terms. Assuming that the set of unitaries \mathbb{U}_L generated by U_L is a 2-design, then

$$\begin{aligned} \mathbb{E}_{\mathbb{U}_L}[f(\theta + \pi \mathbf{e}_i)^2] &= \text{Tr}[\mathbb{E}_{\mathbb{U}_L}[\underbrace{OU_L P_i U_R |0\rangle\langle 0| U_R^\dagger P_i U_L^\dagger}_B \underbrace{OU_L P_i U_R |0\rangle\langle 0| U_R^\dagger P_i U_L^\dagger}_B]] \end{aligned} \tag{B.31}$$

$$= \text{Tr}\left[\frac{\text{Tr}[B^2] \text{Tr}[O]O + \text{Tr}[B]^2 O^2}{2^{2n} - 1} - \frac{\text{Tr}[B^2] O^2 + \text{Tr}[B]^2 \text{Tr}[O]O}{2^n(2^{2n} - 1)}\right] \tag{B.32}$$

$$= \frac{\text{Tr}[O]^2 + \text{Tr}[O^2]}{2^{2n} - 1} - \frac{\text{Tr}[O^2] + \text{Tr}[O]^2}{2^n(2^{2n} - 1)} = \frac{1}{2^n + 1}, \tag{B.33}$$

where in the second line we made use of Eq. (B.22), and the third line the used that $\text{Tr}[B] = \text{Tr}[B^2] = 1$ since $B = P_i U_R |0\rangle\langle 0| U_R^\dagger P_i$ is a projector, and that $\text{Tr}[O] = 0$ and $\text{Tr}[O^2] = 2^n$. Similarly, one can show that integration over \mathbb{U}_R yields the same result. Also, the same calculation leads to $\mathbb{E}_{\mathbb{U}_L}[f(\theta)^2] = \mathbb{E}_{\mathbb{U}_R}[f(\theta)^2] = 1/(2^n + 1)$. Thus, if either \mathbb{U}_L or \mathbb{U}_R is a 2-design then

$$\begin{aligned} \mathbb{E}_{\mathbb{U}_R, \mathbb{U}_L}[f(\theta)^2] &= \mathbb{E}_{\mathbb{U}_R, \mathbb{U}_L}[f(\theta + \pi \mathbf{e}_i)^2] \\ &= \frac{1}{2^n + 1} \quad \forall i \text{ if either } \mathbb{U}_L \text{ or } \mathbb{U}_R \text{ is a 2-design.} \end{aligned} \tag{B.34}$$

Now we evaluate the correlation term $\mathbb{E}[f(\boldsymbol{\theta} + \pi \mathbf{e}_i)f(\boldsymbol{\theta})]$. If \mathbb{U}_L is a 2-design, then

$$\begin{aligned} \mathbb{E}_{U_L}[f(\boldsymbol{\theta} + \pi \mathbf{e}_i)f(\boldsymbol{\theta})] &= \text{Tr}[\mathbb{E}_{U_L}[OU_L P_i U_R |0\rangle\langle 0|U_R^\dagger U_L^\dagger OU_L U_R |0\rangle\langle 0|U_R^\dagger P_i U_L^\dagger]] \\ &= \text{Tr}\left[\frac{\text{Tr}[P_i U_R |0\rangle\langle 0|U_R^\dagger]^2 O^2}{2^{2n} - 1} - \frac{O^2}{2^n(2^{2n} - 1)}\right] \\ &= \frac{1}{2^{2n} - 1} [2^n \text{Tr}[P_i U_R |0\rangle\langle 0|U_R^\dagger]^2 - 1]. \end{aligned} \tag{B.35}$$

While if \mathbb{U}_R is a 2-design instead it holds

$$\begin{aligned} \mathbb{E}_{U_R}[f(\boldsymbol{\theta} + \pi \mathbf{e}_i)f(\boldsymbol{\theta})] &= \text{Tr}[OU_L P_i \mathbb{E}_{U_R}[U_R |0\rangle\langle 0|U_R^\dagger U_L^\dagger OU_L U_R |0\rangle\langle 0|U_R^\dagger] P_i U_L^\dagger] \\ &= \text{Tr}\left[OU_L P_i \frac{(2^n - 1)U_L^\dagger OU_L}{2^n(2^{2n} - 1)} P_i U_L^\dagger\right] \\ &= \frac{1}{2^n(2^n + 1)} \text{Tr}[OU_L P_i U_L^\dagger OU_L P_i U_L^\dagger]. \end{aligned} \tag{B.36}$$

If both of them are 2-designs, then continuing from Eq. (B.36), one obtains

$$\begin{aligned} \mathbb{E}_{U_L, U_R}[f(\boldsymbol{\theta} + \pi \mathbf{e}_i)f(\boldsymbol{\theta})] &= \frac{1}{2^n(2^n + 1)} \text{Tr}[\mathbb{E}_{U_L}[OU_L P_i U_L^\dagger OU_L P_i U_L^\dagger]] \\ &= \frac{1}{2^n(2^n + 1)} \text{Tr}\left[\frac{\text{Tr}[P_i]^2 O^2 + \text{Tr}[P_i^2] \text{Tr}[O]O}{2^{2n} - 1} - \frac{\text{Tr}[P_i^2] O^2 + \text{Tr}[P_i] \text{Tr}[O]O}{2^n(2^{2n} - 1)}\right] \\ &= -\frac{1}{2^n(2^n + 1)} \frac{\text{Tr}[P_i^2] \text{Tr}[O^2]}{2^n(2^{2n} - 1)} = -\frac{1}{(2^n + 1)(2^{2n} - 1)} \in \mathcal{O}(2^{-3n}). \end{aligned} \tag{B.37}$$

Finally, plugging Eqs. (B.35), (B.36) and (B.37) in Eq. (B.29), one has $\forall i = 1, \dots, M$

$$\begin{aligned} \text{Var}_{U_L, U_R}[H_{ii}] &= \frac{1}{2} \mathbb{E}[f(\boldsymbol{\theta})^2] - \frac{1}{2} \mathbb{E}[f(\boldsymbol{\theta} + \pi \mathbf{e}_i)f(\boldsymbol{\theta})] \\ &= \frac{1}{2(2^n + 1)} - \frac{1}{2} \begin{cases} \frac{1}{2^{2n} - 1} [2^n \text{Tr}[P_i U_R |0\rangle\langle 0|U_R^\dagger]^2 - 1] & \forall i, \text{ if } \mathbb{U}_L \text{ 2-design,} \\ \frac{1}{2^n(2^n + 1)} \text{Tr}[OU_L P_i U_L^\dagger OU_L P_i U_L^\dagger] & \forall i, \text{ if } \mathbb{U}_R \text{ 2-design,} \\ -\frac{1}{(2^n + 1)(2^{2n} - 1)} & \forall i, \text{ if } \mathbb{U}_L, \mathbb{U}_R \text{ 2-designs,} \end{cases} \end{aligned} \tag{B.38}$$

where $\mathbb{U}_R = \mathbb{U}_R^{(i)}$ and $\mathbb{U}_L = \mathbb{U}_L^{(i)}$ are defined as in Eq. (B.24) and actually depend on the index i of the parameter.

Not surprisingly, as it happens for first order derivatives, also second order derivatives of PQCs are found to be exponentially vanishing [36, 56], as from Eq. (B.38) one can check that $\text{Var}[H_{ii}] \in \mathcal{O}(2^{-n})$.

B.2.2 Statistics of the trace of the Hessian

The average value of the trace of the Hessian is easily found to be zero using Eq. (B.28), in fact

$$\mathbb{E}_{U_R, U_L}[\text{Tr}[H]] = \sum_{i=1}^M \mathbb{E}_{U_R^{(i)}, U_L^{(i)}}[H_{ii}] = 0, \tag{B.39}$$

where we assume that for every parameter i either $U_R^{(i)}$ or $U_L^{(i)}$ is a 1-design. The variance of the trace is instead

$$\text{Var}_{U_R, U_L}[\text{Tr}[H]] = \text{Var}\left[\sum_{i=1}^M H_{ii}\right] = \sum_{i=1}^M \text{Var}[H_{ii}] + 2 \sum_{i<j}^M \text{Cov}[H_{ii}H_{jj}]. \tag{B.40}$$

We can upper bound this quantity using the covariance inequality [79]

$$|\text{Cov}[H_{ii}, H_{jj}]| \leq \sqrt{\text{Var}[H_{ii}] \text{Var}[H_{jj}]} \approx \text{Var}[H_{ii}],$$

were we assumed that $\text{Var}[H_{ii}] \approx \text{Var}[H_{jj}] \forall i, j$. Using that $\text{Var}[H_{ii}] \in \mathcal{O}(2^{-n})$ one finally has

$$\text{Var}_{U_R, U_L}[\text{Tr}[H]] \leq \sum_{i=1}^M \text{Var}[H_{ii}] + 2 \sum_{i<j}^M \text{Var}[H_{ii}] \in \mathcal{O}\left(\frac{M^2}{2^n}\right). \tag{B.41}$$

Alternatively, one can obtain a tighter yet qualitative approximation by explicitly considering the nature of the sums in Eq. (B.40). First, by using Eq. (B.23), the covariance term is explicitly

$$\begin{aligned} \text{Cov}[H_{ii}, H_{jj}] &= \mathbb{E}[H_{ii}H_{jj}] = \frac{1}{4} \mathbb{E}[(f_i - f)(f_j - f)] \\ &= \frac{1}{4} \mathbb{E}[f^2] + \frac{1}{4} \mathbb{E}[f_i f_j] - \frac{1}{4} \mathbb{E}[f_i f] - \frac{1}{4} \mathbb{E}[f_j f], \end{aligned} \tag{B.42}$$

where for ease of notation we defined $f_{i,j} = f(\theta + \pi \mathbf{e}_{i,j})$ and $f = f(\theta)$. Note that except for the first term which is always positive, all remaining correlations terms can be both positive and negative. Also, all of these terms are bounded from above by the same quantity, as via Cauchy-Schwarz it follows

$$|\mathbb{E}[f_i f_j]| \leq \sqrt{\mathbb{E}[f_i^2] \mathbb{E}[f_j^2]} = \frac{1}{2^n + 1} \quad \text{and} \quad |\mathbb{E}[f_i f]| \leq \sqrt{\mathbb{E}[f_i^2] \mathbb{E}[f^2]} = \frac{1}{2^n + 1}, \tag{B.43}$$

where we have used $E[f^2] = E[f_i^2] = 1/(2^n + 1)$ from Eq. (B.34). Then, the variance can be written as

$$\begin{aligned} &\text{Var}_{U_R, U_L}[\text{Tr}[H]] \\ &= \sum_{i=1}^M \text{Var}[H_{ii}] + 2 \sum_{i<j}^M \mathbb{E}[H_{ii}H_{jj}] \\ &= \sum_{i=1}^M \frac{\mathbb{E}[f^2] - \mathbb{E}[f_i f]}{2} + 2 \sum_{i<j}^M \frac{\mathbb{E}[f^2] + \mathbb{E}[f_i f_j] - \mathbb{E}[f_i f] - \mathbb{E}[f_j f]}{4} \end{aligned}$$

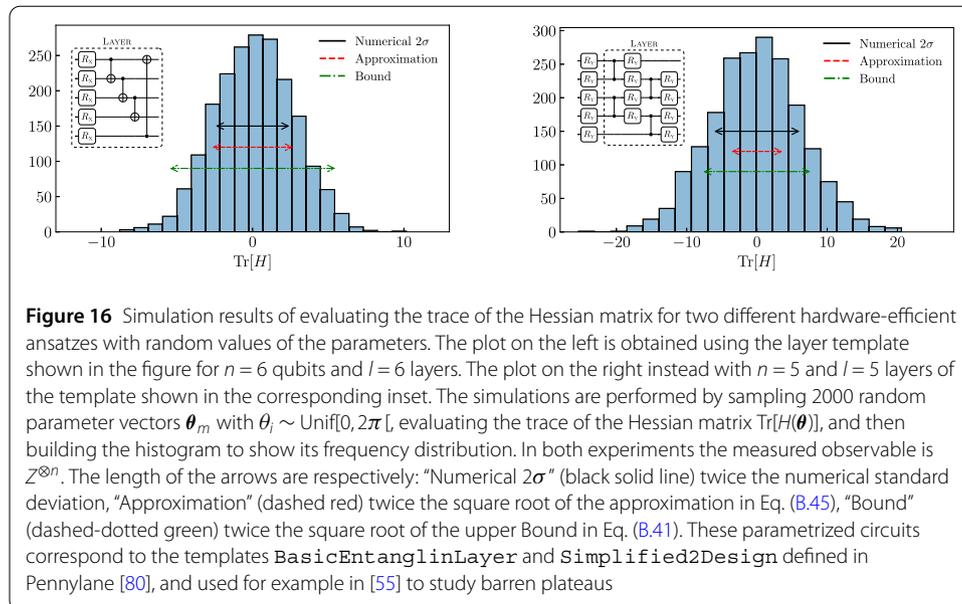
$$\begin{aligned}
 &= \frac{1}{2} \left(\sum_{i=1}^M + \sum_{i<j}^M \right) \mathbb{E}[f^2] - \frac{1}{2} \left(\sum_{i=1}^M \mathbb{E}[f_i f] + \sum_{i<j}^M \mathbb{E}[f_i f] + \sum_{i<j}^M \mathbb{E}[f_j f] \right) + \frac{1}{2} \sum_{i<j}^M \mathbb{E}[f_i f_j] \\
 &= \frac{M(M+1)}{4} \mathbb{E}[f^2] - \underbrace{\frac{M}{2} \sum_{i=1}^M \mathbb{E}[f_i f] + \frac{1}{2} \sum_{i<j}^M \mathbb{E}[f_i f_j]}_{\Delta}. \tag{B.44}
 \end{aligned}$$

Numerical simulations In addition to Fig. 6 in the main text, in Fig. 16 we report numerical evidence for the trace of the Hessian for two common hardware-efficient parametrized quantum circuit ansatzes. The histograms represent the frequency of obtaining a given value of the trace of the Hessian $\text{Tr}[H(\theta)]$ upon random assignments of the parameters. The length of the arrows are, respectively: “Numerical 2σ ” (black solid line) twice the statistical standard deviation computed from the numerical results, “Approximation” (dashed red) twice the square root of the Eq. (B.44) with $\Delta = 0$, “Bound” (dashed-dotted green) twice the square root of the upper Bound in Eq. (B.41).

All simulations confirm the bound (B.41), and, more interestingly, both the circuit on the left of Fig. 16 and the one in Fig. 6 in the main text, have a numerical variance which is very well approximated by Eq. (B.44) with $\Delta = 0$. We conjecture this is due to the fact that all correlation terms in Eq. (B.44) are roughly of the same order of magnitude (see Eq. (B.43)), and can be either positive and negative, depending on the parameter and the specifics of the ansatz. Thus, one can expect the whole contribution to either vanish $\Delta \approx 0$, or be negligible with respect to the leading term. If this is the case, then substituting $\mathbb{E}[f^2] = 1/(2^n + 1)$, the variance of the Hessian is approximately

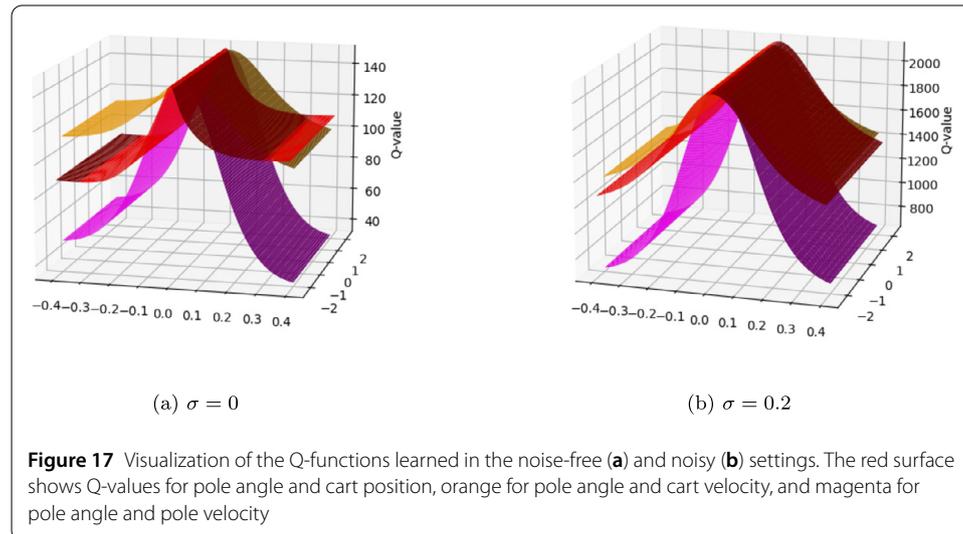
$$\text{Var}_{U_R, U_L}[\text{Tr}[H]] \approx \frac{M(M+1)}{4} \mathbb{E}[f^2] = \frac{M(M+1)}{4(2^n + 1)} \approx \frac{1}{4} \frac{M^2}{2^n}, \tag{B.45}$$

which is four times smaller than the upper bound Eq. (B.41), but clearly has the same scaling. While we numerically verified it also at other number of qubits, more investigations



are needed to understand if and when this approximation holds, and we leave a detailed study of this phenomenon for future work.

Appendix C: Visualization of CartPole policies obtained with Q-learning



Acknowledgements

AS is funded by the German Ministry for Education and Research (BMB+F) in the project QAI2-Q-KIS under grant 13N15587. This work was also supported by the Dutch Research Council (NWO/OCW), as part of the Quantum Software Consortium programme (project number 024.003.037). CM acknowledges support by the National Research Centre for HPC, Big Data and Quantum Computing (ICSC: MUR project CN0000013).

Availability of data and materials

The code that was used to generate the numerical results in this work can be found on GitHub (https://github.com/askolik/noisy_qrl), along with the data set containing the TSP instances studied in this work and their optimal solutions.

Declarations

Competing interests

The authors declare no competing interests.

Author contributions

AS conceived the idea for this work and conducted the numerical experiments. SM performed analytical study on the effect of Gaussian noise and provided decoherence noise model. AS and VD proposed shot allocation algorithm. AS and SM wrote the first version of the manuscript, all authors contributed to the final editing.

Author details

¹Leiden University, Niels Bohrweg 1, 2333 CA Leiden, The Netherlands. ²Volkswagen AG, Ungererstraße 69, 80805 Munich, Germany. ³Dipartimento di Fisica, Università di Pavia, Via Bassi 6, I-27100, Pavia, Italy. ⁴INFN Sezione di Pavia, Via Bassi 6, I-27100, Pavia, Italy. ⁵CNR-INO, Largo E. Fermi 6, I-50125, Firenze, Italy.

Publisher's Note

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

Received: 16 January 2023 Accepted: 13 February 2023 Published online: 28 February 2023

References

1. Bharti K, Cervera-Lierta A, Kyaw TH, Haug T, Alperin-Lea S, Anand A, Degroote M, Heimonen H, Kottmann JS, Menke T, Mok W-K, Sim S, Kwek L-C, Aspuru-Guzik A. Noisy intermediate-scale quantum algorithms. *Rev Mod Phys.* 2022;94:015004.
2. Cerezo M, Arrasmith A, Babbush R, Benjamin SC, Endo S, Fujii K, McClean JR, Mitarai K, Yuan X, Cincio L et al. Variational quantum algorithms. *Nat Rev Phys.* 2021;3(9):625–44.
3. Mangini S, Tacchino F, Gerace D, Bajoni D, Macchiavello C. Quantum computing models for artificial neural networks. *Europhys Lett.* 2021;134(1):10002.
4. Gentini L, Cuccoli A, Pirandola S, Verrucchi P, Banchi L. Noise-resilient variational hybrid quantum-classical optimization. *Phys Rev A.* 2020;102:052414.
5. Jim K-C, Giles CL, Horne BG. An analysis of noise in recurrent neural networks: convergence and generalization. *IEEE Trans Neural Netw.* 1996;7(6):1424–38.
6. Noh H, You T, Mun J, Han B. Regularizing deep neural networks by noise: its interpretation and optimization. In: *Advances in neural information processing systems.* vol. 30. 2017.
7. Graves A. Practical variational inference for neural networks. In: *Advances in neural information processing systems.* vol. 24. 2011.
8. Graves A, Mohamed A, Hinton G. Speech recognition with deep recurrent neural networks. In: 2013 IEEE international conference on acoustics, speech and signal processing. *Ieee;* 2013. p. 6645–9.
9. Balda ER, Behboodi A, Mathar R. Adversarial examples in deep neural networks: an overview. In: *Deep learning: algorithms and applications;* 2020. p. 31–65.
10. Xie C, Wang J, Zhang Z, Ren Z, Yuille A. Mitigating adversarial effects through randomization. 2017. arXiv preprint. [arXiv:1711.01991](https://arxiv.org/abs/1711.01991).
11. Gilmer J, Ford N, Carlini N, Cubuk E. Adversarial examples are a natural consequence of test error in noise. In: *International conference on machine learning.* PMLR; 2019. p. 2280–9.
12. Jaekle F, Kumar MP. Generating adversarial examples with graph neural networks. In: *Uncertainty in artificial intelligence.* PMLR; 2021. p. 1556–64.
13. Goodfellow IJ, Shlens J, Szegedy C. Explaining and harnessing adversarial examples. 2014. arXiv preprint. [arXiv:1412.6572](https://arxiv.org/abs/1412.6572).
14. Srivastava N, Hinton G, Krizhevsky A, Sutskever I, Salakhutdinov R. Dropout: a simple way to prevent neural networks from overfitting. *J Mach Learn Res.* 2014;15(1):1929–58.
15. Wang S, Fontana E, Cerezo M, Sharma K, Sone A, Cincio L, Coles PJ. Noise-induced barren plateaus in variational quantum algorithms. *Nat Commun.* 2021;12(1):6961.
16. Zeng J, Wu Z, Cao C, Zhang C, Hou S-Y, Xu P, Zeng B. Simulating noisy variational quantum eigensolver with local noise models. *Quantum Eng.* 2021;3(4):e77.
17. Farhi E, Goldstone J, Gutmann S. A quantum approximate optimization algorithm. 2014. arXiv preprint. [arXiv:1411.4028](https://arxiv.org/abs/1411.4028).
18. Alam M, Ash-Saki A, Ghosh S. Analysis of quantum approximate optimization algorithm under realistic noise in superconducting qubits. 2019. arXiv preprint. [arXiv:1907.09631](https://arxiv.org/abs/1907.09631).
19. Harrigan MP, Sung KJ, Neeley M, Satzinger KJ, Arute F, Arya K, Atalaya J, Bardin JC, Barends R, Boixo S et al. Quantum approximate optimization of non-planar graph problems on a planar superconducting processor. *Nat Phys.* 2021;17(3):332–6.
20. LaRose R, Coyle B. Robust data encodings for quantum classifiers. *Phys Rev A.* 2020;102:032420.
21. Liu J, Wilde F, Mele AA, Jiang L, Eisert J. Noise can be helpful for variational quantum algorithms. 2022. arXiv preprint. [arXiv:2210.06723](https://arxiv.org/abs/2210.06723).
22. Wang J, Liu Y, Li B. Reinforcement learning with perturbed rewards. In: *Proceedings of the AAAI conference on artificial intelligence.* vol. 34. 2020. p. 6202–9.
23. Huang S, Papernot N, Goodfellow I, Duan Y, Abbeel P. Adversarial attacks on neural network policies. 2017. arXiv preprint. [arXiv:1702.02284](https://arxiv.org/abs/1702.02284).
24. Kos J, Song D. Delving into adversarial attacks on deep policies. 2017. arXiv preprint. [arXiv:1705.06452](https://arxiv.org/abs/1705.06452).
25. Yu Y. Towards sample efficient reinforcement learning. In: *IJCAI.* 2018. p. 5739–43.
26. Chen SY-C, Yang C-HH, Qi J, Chen P-Y, Ma X, Goan H-S. Variational quantum circuits for deep reinforcement learning. *IEEE Access.* 2020;8:141007–24.
27. Lockwood O, Si M. Reinforcement learning with quantum variational circuit. In: *Proceedings of the AAAI conference on artificial intelligence and interactive digital entertainment.* vol. 16. 2020. p. 245–51.
28. Jerbi S, Gyurik C, Marshall S, Briegel H, Dunjko V. Parametrized quantum policies for reinforcement learning. *Adv Neural Inf Process Syst.* 2021;34:28362–75.
29. Skolik A, Jerbi S, Dunjko V. Quantum agents in the gym: a variational quantum algorithm for deep q-learning. *Quantum.* 2022;6:720.
30. Lan Q. Variational quantum soft actor-critic. 2021. arXiv preprint. [arXiv:2112.11921](https://arxiv.org/abs/2112.11921).
31. Wu S, Jin S, Wen D, Wang X. Quantum reinforcement learning in continuous action space. 2020. arXiv preprint. [arXiv:2012.10711](https://arxiv.org/abs/2012.10711).
32. Sequeira A, Santos LP, Barbosa LS. Variational quantum policy gradients with an application to quantum control. 2022. arXiv preprint. [arXiv:2203.10591](https://arxiv.org/abs/2203.10591).
33. Lockwood O, Si M. Playing atari with hybrid quantum-classical reinforcement learning. In: *NeurIPS 2020 workshop on pre-registration in machine learning.* PMLR; 2021. p. 285–301.
34. Franz M, Wolf L, Periyasamy M, Ufrecht C, Scherer DD, Plinge A, Mutschler C, Mauerer W. Uncovering instabilities in variational-quantum deep q-networks. 2022. arXiv preprint. [arXiv:2202.05195](https://arxiv.org/abs/2202.05195).
35. Ito K, Mizukami W, Fujii K. Universal noise-precision relations in variational quantum algorithms. 2021. arXiv preprint. [arXiv:2106.03390](https://arxiv.org/abs/2106.03390).
36. Cerezo M, Coles PJ. Higher order derivatives of quantum neural networks with barren plateaus. *Quantum Sci Technol.* 2021;6(3):035006.
37. Sutton RS, Barto AG. *Reinforcement learning: an introduction.* Cambridge: MIT Press; 2018.

38. Mnih V, Kavukcuoglu K, Silver D, Rusu AA, Veness J, Bellemare MG, Graves A, Riedmiller M, Fidjeland AK, Ostrovski G et al. Human-level control through deep reinforcement learning. *Nature*. 2015;518(7540):529–33.
39. Skolik A, Cattelan M, Yarkoni S, Bäck T, Dunjko V. Equivariant quantum circuits for learning on weighted graphs. 2022. arXiv preprint. [arXiv:2205.06109](https://arxiv.org/abs/2205.06109).
40. Skolik A, Mangini S. Code that was used for training of noisy quantum agents. 2022. https://github.com/askolik/noisy_qrl.
41. Openai gym. <https://github.com/openai/gym/wiki>. Accessed: 06-09-2022.
42. Lillicrap TP, Hunt JJ, Pritzel A, Heess N, Erez T, Tassa Y, Silver D, Wierstra D. Continuous control with deep reinforcement learning. 2015. arXiv preprint. [arXiv:1509.02971](https://arxiv.org/abs/1509.02971).
43. Kandala A, Mezzacapo A, Temme K, Takita M, Brink M, Chow JM, Gambetta JM. Hardware-efficient variational quantum eigensolver for small molecules and quantum magnets. *Nature*. 2017;549(7671):242–6.
44. Pérez-Salinas A, Cervera-Lierta A, Gil-Fuster E, Latorre JI. Data re-uploading for a universal quantum classifier. *Quantum*. 2020;4:226.
45. Schuld M, Sweke R, Meyer JJ. Effect of data encoding on the expressive power of variational quantum-machine-learning models. *Phys Rev A*. 2021;103:032430.
46. Tensorflow quantum rl tutorial. https://www.tensorflow.org/quantum/tutorials/quantum_reinforcement_learning. Accessed: 06-09-2022.
47. Bello I, Pham H, Le QV, Norouzi M, Bengio S. Neural combinatorial optimization with reinforcement learning. 2016. arXiv preprint. [arXiv:1611.09940](https://arxiv.org/abs/1611.09940).
48. Slivkins A et al. Introduction to multi-armed bandits. *Found Trends Mach Learn*. 2019;12(1–2):1–286.
49. Lai TL, Robbins H et al. Asymptotically efficient adaptive allocation rules. *Adv Appl Math*. 1985;6(1):4–22.
50. Auer P. Using confidence bounds for exploitation-exploration trade-offs. *J Mach Learn Res*. 2002;3(Nov):397–422.
51. Cai Z, Xu X, Benjamin SC. Mitigating coherent noise using pauli conjugation. *npj Quantum Inf*. 2020;6(1):1–9.
52. Schulm M, Bergholm V, Gogolin C, Izaac J, Killoran N. Evaluating analytic gradients on quantum hardware. *Phys Rev A*. 2019;99:032331.
53. Mitarai K, Negoro M, Kitagawa M, Fujii K. Quantum circuit learning. *Phys Rev A*. 2018;98:032309.
54. McClean JR, Boixo S, Smelyanskiy VN, Babbush R, Neven H. Barren plateaus in quantum neural network training landscapes. *Nat Commun*. 2018;9(1):4812.
55. Cerezo M, Sone A, Volkoff T, Cincio L, Coles PJ. Cost function dependent barren plateaus in shallow parametrized quantum circuits. *Nat Commun*. 2021;12(1):1791.
56. Holmes Z, Sharma K, Cerezo M, Coles PJ. Connecting ansatz expressibility to gradient magnitudes and barren plateaus. *PRX Quantum*. 2022;3:010313.
57. Huang H-Y, Kueng R, Preskill J. Predicting many properties of a quantum system from very few measurements. *Nat Phys*. 2020;16(10):1050–7.
58. Puchała Z, Miszczyk JA. Symbolic integration with respect to the Haar measure on the unitary groups. *Bull Pol Acad Sci, Tech Sci*. 2017;65(1):21–7.
59. Van Hasselt H, Guez A, Silver D. Deep reinforcement learning with double q-learning. In: *Proceedings of the AAAI conference on artificial intelligence*. vol. 30. 2016.
60. Broughton M, Verdon G, McCourt T, Martinez AJ, Yoo JH, Isakov SV, Massey P, Halavati R, Niu MY, Zlokapa A, et al. Tensorflow quantum: a software framework for quantum machine learning. 2020. arXiv preprint. [arXiv:2003.02989](https://arxiv.org/abs/2003.02989).
61. Google Inc. Documentation of depolarizing channel in cirq. 2022. <https://quantumai.google/reference/python/cirq/depolarize>.
62. Isakov SV, Kafri D, Martin O, Heidweiller CV, Mruczkiewicz W, Harrigan MP, Rubin NC, Thomson R, Broughton M, Kissell K, Peters E, Gustafson E, Li ACY, Lamm H, Perdue G, Ho AK, Strain D, Boixo S. Simulations of quantum circuits with approximate noise using qsim and cirq. 2021.
63. Nielsen MA, Chuang IL. *Quantum computation and quantum information*. Cambridge: Cambridge University Press; 2010.
64. Proctor T, Seritan S, Rudinger K, Nielsen E, Blume-Kohout R, Young K. Scalable randomized benchmarking of quantum computers using mirror circuits. *Phys Rev Lett*. 2022;129:150502.
65. Vovrosh J, Khosla KE, Greenaway S, Self C, Kim MS, Knolle J. Simple mitigation of global depolarizing errors in quantum simulations. *Phys Rev E*. 2021;104:035309.
66. Magesan E, Gambetta JM, Emerson J. Characterizing quantum gates via randomized benchmarking. *Phys Rev A*. 2012;85:042311.
67. McKay DC, Sheldon S, Smolin JA, Chow JM, Gambetta JM. Three-qubit randomized benchmarking. *Phys Rev Lett*. 2019;122:200502.
68. Ryan-Anderson C, Brown NC, Allman MS, Arkin B, Asa-Attuah G, Baldwin C, Berg J, Bohnet JG, Braxton S, Burdick N, Campora JP, Chernoguzov A, Esposito J, Evans B, Francois D, Gaebler JP, Gatterman TM, Gerber J, Gilmore K, Gresh D, Hall A, Hankin A, Hostetter J, Lucchetti D, Mayer K, Myers J, Neyenhuis B, Santiago J, Sedlacek J, Skripka T, Slattery A, Stutz RP, Tait J, Tobey R, Vittorini G, Walker J, Hayes D. 2022.
69. [ibmqquantum](https://quantum-computing.ibm.com/). 2022. <https://quantum-computing.ibm.com/>.
70. Pelofske E, Bäertschi A, Eidenbenz S. Quantum volume in practice: what users can expect from nisq devices. 2022. arXiv preprint. [arXiv:2203.03816](https://arxiv.org/abs/2203.03816).
71. IBM Quantum Experience. <https://quantum-computing.ibm.com/services/resources?tab=systems>; 2022.
72. França DS, Garcia-Patron R. Limitations of optimization algorithms on noisy quantum devices. *Nat Phys*. 2021;17(11):1221–7.
73. Gao X, Duan L. Efficient classical simulation of noisy quantum computation. 2018. arXiv preprint. [arXiv:1810.03176](https://arxiv.org/abs/1810.03176).
74. LaRose R, Mari A, Kaiser S, Karalekas PJ, Alves AA, Czarnik P, El Mandouh M, Gordon MH, Hindy Y, Robertson A, Thakre P, Wahl M, Samuel D, Mistri R, Tremblay M, Gardner N, Stemen NT, Shammah N, Zeng WJ. Mitiq: a software package for error mitigation on noisy quantum computers. *Quantum*. 2022;6:774.
75. Russo V, Mari A, Shammah N, LaRose R, Zeng WJ. Testing platform-independent quantum error mitigation on noisy quantum computers. 2022.
76. Wang S, Czarnik P, Arrasmith A, Cerezo M, Cincio L, Coles PJ. Can error mitigation improve trainability of noisy variational quantum algorithms? 2021.

77. Huembeli P, Dauphin A. Characterizing the loss landscape of variational quantum circuits. *Quantum Sci Technol*. 2021;6(2):025011.
78. Fukuda M, König R, Nechita I. RTNI—a symbolic integrator for Haar-random tensor networks. *J Phys A, Math Theor*. 2019;52(42):425303.
79. Keener RW. *Theoretical statistics: topics for a core course*. 1st ed. Springer texts in statistics. Berlin: Springer; 2010.
80. Bergholm V, Izaac J, Schuld M, Gogolin C, Ahmed S, Ajith V, Alam MS, Alonso-Linaje G, AkashNarayanan B, Asadi A, Arrazola JM, Azad U, Banning S, Blank C, Bromley TR, Cordier BA, Ceroni J, Delgado A, Di Matteo O, Dusko A, Garg T, Guala D, Hayes A, Hill R, Ijaz A, Isacsson T, Ittah D, Jahangiri S, Jain P, Jiang E, Khandelwal A, Kottmann K, Lang RA, Lee C, Loke T, Lowe A, McKiernan K, Meyer JJ, Montañez-Barrera JA, Moyard R, Niu Z, O’Riordan LJ, Oud S, Panigrahi A, Park C-Y, Polatajko D, Quesada N, Roberts C, Sá N, Schoch I, Shi B, Shu S, Sim S, Singh A, Strandberg I, Soni J, Száva A, Thabet S, Vargas-Hernández RA, Vincent T, Vitucci N, Weber M, Wierichs D, Wiersema R, Willmann M, Wong V, Zhang S, Killoran N. *PennyLane: automatic differentiation of hybrid quantum-classical computations*. 2018.

Submit your manuscript to a SpringerOpen[®] journal and benefit from:

- ▶ Convenient online submission
- ▶ Rigorous peer review
- ▶ Open access: articles freely available online
- ▶ High visibility within the field
- ▶ Retaining the copyright to your article

Submit your next manuscript at ▶ [springeropen.com](https://www.springeropen.com)
