Quantum machine learning: on the design, trainability and noise-robustness of near-term algorithms

Skolik, A.

Appendix

Layerwise learning for quantum neural networks

As alluded to in Section 4.1, LL and CDL perform similarly in a perfect simulation scenario, where we assume neither shot nor hardware noise. Figure 1 a) shows a comparison of LL and CDL under perfect conditions, i.e. infinite number of measurements and a batch size that corresponds to the number of samples, which enables computation of exact gradients. Here, the magnitude of gradients doesn't affect the learning process severely, as the Adam optimizer uses adaptive learning rates for each parameter and can therefore handle different ranges of gradient magnitudes well as long as there is some variance in the computed gradients. In this regime, both approaches show similar performance.

The convergence rate of a PQC increases proportionally to the number of parameters in a model [188], so the number of experiment repetitions is almost equal for LL and CDL. LL has less parameters and needs more epochs to converge due to this, whereas CDL needs more calls to the quantum device for one update step, but in turn needs less epochs to converge. In terms of cross entropy, both LL and CDL converge to a value of roughly 0.51. The corresponding test error of all approaches, except for the randomly initialized CDL, reaches almost 0 but doesn't converge there and settles around an error of roughly 0.1 eventually, as seen in Figure 1 b).

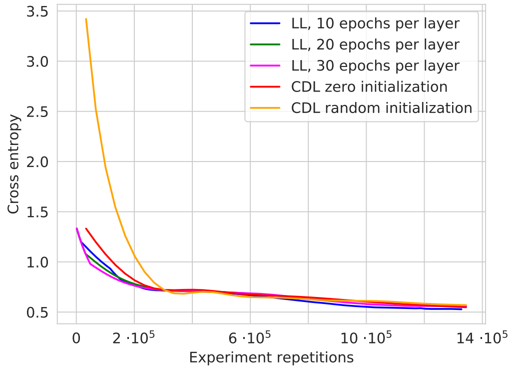

(a) cross entropy

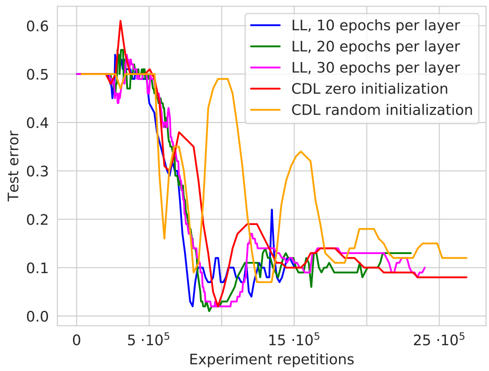

(b) test error

Figure 1: a) Cross entropy of LL and CDL during training with exact gradient calculation corresponding to infinite number of measurements. When one assumes the unphysical situation of infinite measurements ($m = \infty$) all methods seem to perform similarly. In particular, we compare LL to CDL with zero and random initialization, where the initial parameters for the latter are chosen uniformly from $[0, 2\pi)$. The hyperparameters for all configurations were set to $m = \infty$, $b = 100$ and $\eta = 0.01$. (For computing the number of experiment repetitions as defined in Section 4.2.3, we drop m.) b) Test error corresponding to the runs shown in Figure 1. This further supports the observation that when one allows unphysical, arbitrary precision queries($m = \infty$), all tuned training strategies seem to perform similarly.

Quantum agents in the Gym: A variational quantum algorithm for deep Q-learning

Visualization of a learned Q-function

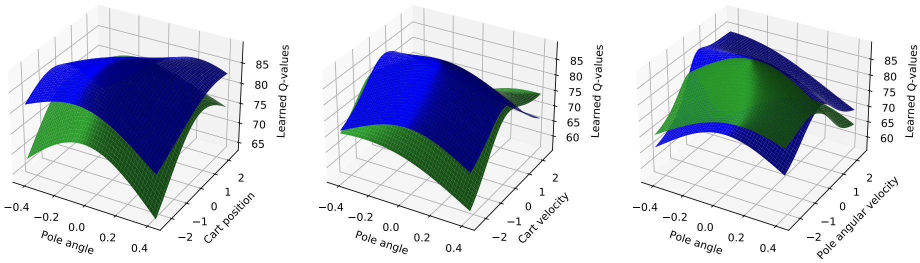

Figure 2: Visualization of the approximate Q-function learned by a quantum Q-learning agent solving Cart Pole. Due to the 4 dimensions of the state space in Cart Pole, we represent the Q-values associated to the actions "left" (green) and "right" (blue) on 3 subspaces of the state space by fixing unrepresented dimensions to 0 in each plot. As opposed to the analogue values (i.e., unnormalized policy) learned by policy-gradient PQC agents in this environment [150], the approximate Q-values appear nicely-behaved, likely due to the stronger constraints that Q-learning has on well-performing function approximations.

Model hyperparameters

In the following, we give a detailed list of the hyperparameters for each configuration in fig. 5.3, fig. 5.4, fig. 5.5, fig. 5.6 and fig. 5.7. The hyperparameters that we searched over for each model were the following (see explanations of each hyperparameter in table 1):

- *Frozen Lake v0*: update model, update target model, η
- *Cart Pole v0, quantum model*: batch size, update model, update target model, η, train w_d, train w_o, η_{w_d}, η_{w_o}
- *Cart Pole v0, classical model*: number of units per layer, batch size, update model, update target model, η

	Hyperparameter explanation
qubits	number of qubits in circuit
layers	number of layers
γ	discount factor for Q-learning
train w_d	train weights on the model input as defined in section 5.1.1
train w_o	train weights on the model output as defined in section 5.1.2
η	model parameter learning rate
η_{w_d}	input weight learning rate
η_{w_o}	output weight learning rate
batch size	number of samples shown to optimizer at each update
ϵ_{init}	initial value for ϵ-greedy policy
ϵ_{dec}	decay of ϵ for ϵ-greedy policy
ϵ_{min}	minimal value of ϵ for ϵ-greedy policy
update model	time steps after which model is updated
update target model	time steps after which model parameters are copied to target model
size of replay memory	size of memory for experience replay
data re-uploading	use data re-uploading as defined in section 5.1.1

Table 1: Description of hyperparameters considered in this work

	Frozen Lake v0, fig. 5.3	**Cart Pole v0, optimal**	**Cart Pole v0, sub-optimal**
qubits	4	4	4
layers	5, 10, 15	5	5
γ	0.8	0.99	0.99
train w_d	no	yes, no	yes, no
train w_o	no	yes, no	yes, no
η	0.001	0.001	0.001
η_{w_d}	–	0.001	0.001
η_{w_o}	–	0.1	0.1
batch size	11	16	16
ϵ_{init}	1	1	1
ϵ_{dec}	0.99	0.99	0.99
ϵ_{min}	0.01	0.01	0.01
update model	5	1	10
update target model	10	1	30
size of replay memory	10000	10000	10000
data re-uploading	no	yes, no	yes, no

Table 2: Hyperparameter settings of PQCs in fig. 5.3, fig. 5.4 and fig. 5.5

layers	5	10	15	20	25	30
qubits	4	4	4	4	4	4
γ	0.99	0.99	0.99	0.99	0.99	0.99
train w_d	yes	yes	yes	yes	yes	yes
train w_o	yes	yes	yes	yes	yes	yes
η	0.001	0.001	0.001	0.001	0.001	0.001
η_{w_d}	0.001	0.001	0.001	0.001	0.001	0.001
η_{w_o}	0.1	0.1	0.1	0.1	0.1	0.1
batch size	16	64	32	16	64	16
ϵ_{init}	1	1	1	1	1	1
ϵ_{dec}	0.99	0.99	0.99	0.99	0.99	0.99
ϵ_{min}	0.01	0.01	0.01	0.01	0.01	0.01
update model	1	10	10	10	10	10
update target model	1	30	30	30	30	30
size of replay memory	10000	10000	10000	10000	10000	10000
data re-uploading	yes	yes	yes	yes	yes	yes

Table 3: Hyperparameter settings of PQCs in fig. 5.6 a)

units in hidden layers	(10, 10)	(15, 15)	(20, 20)	(24, 24)	(30, 30)	(64, 64)
γ	0.99	0.99	0.99	0.99	0.99	0.99
η	0.001	0.001	0.001	0.001	0.001	0.001
batch size	64	16	64	64	64	16
ϵ_{init}	1	1	1	1	1	1
ϵ_{dec}	0.99	0.99	0.99	0.99	0.99	0.99
ϵ_{min}	0.01	0.01	0.01	0.01	0.01	0.01
update model	1	1	1	1	1	1
update target model	1	1	1	1	1	1
size of replay memory	10000	10000	10000	10000	10000	10000

Table 4: Hyperparameter settings of NNs in fig. 5.6 b)

Equivariant quantum circuits for learning on weighted graphs

Additional results on statistical significance of comparison between EQC and NEQC

To make statements on the statistical significance of the difference between the performance of the EQC and NEQC shown in Figure 6.5, we perform a two-sample t-test on the two models for the same instance sizes (i.e., for the data in the two boxes for each instance size) with the null hypothesis that the averages of the two distributions are the same. Based on this, we compute p-values to quantify the statistical significance of the differences between models.

Figure 3 a) shows p-values for the depth-one EQCs and NEQCs from Figure 6.5 b). For the 5-city instances, we can not reject the null hypothesis. Indeed, it is already visible by looking at the boxes that the distributions are very similar, which can be expected as the number of permutations of a graph with five vertices is small. However, as we scale up the instance size to ten cities, the corresponding p-value is much smaller than 0.05, which means that we can reject the null hypothesis that the two distributions have the same average with high confidence. This is also the case for the instances with twenty cities, where the p-value is less than 0.01.

Figure 3 b) shows p-values for the depth four EQCs and NEQCs from Figure 6.5 d). Again, the p-value of the 5-city instances is very high with 0.74, so that we can not reject the null hypothesis. Also similarly to the above, the p-values get smaller as we scale up the instance size. For the depth-four ansatzes, the p-value is smallest for the twenty city instances, with a value much smaller than 0.05.

To provide additional insight, we also plot the means and their standard error for both the 1-layer (EQC-1, NEQC-1) and 4-layer (EQC-4, NEQC-4) models in Figure 6.5. As a rule of thumb, one can expect that when the error bars given by the standard errors of two means do not overlap, the p-value can be smaller than 0.05, while in the case that they do overlap, the p-value is likely much larger. The error bars in Figure 4 are in line with this statement, where we see that the error bars for the five-city instances overlap for both circuit depths, while this is not the case for the larger instance sizes and in addition the distance between the means increases for those instance sizes. Remarkably, we also see that the difference

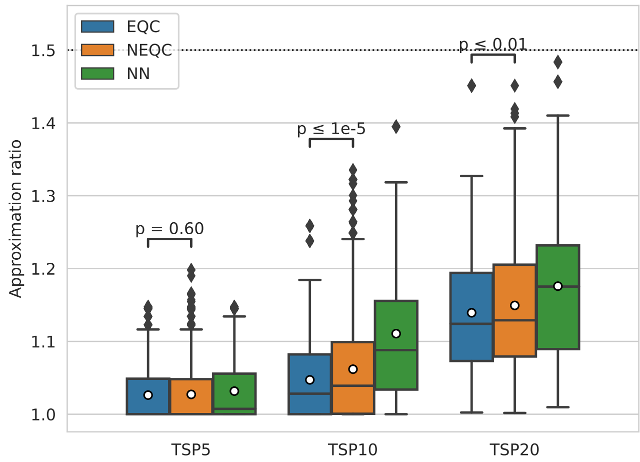

(a) one layer

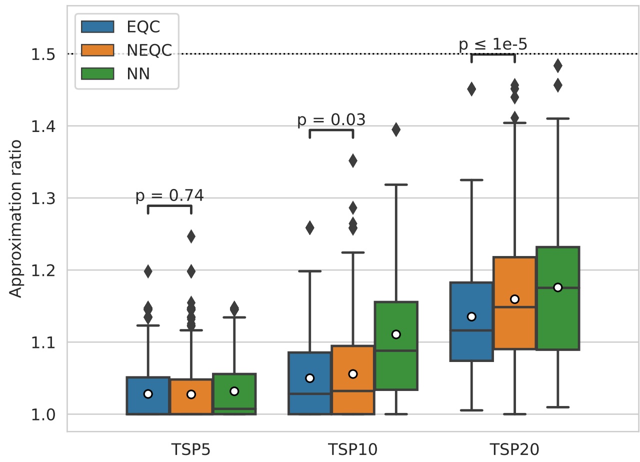

(b) four layers

Figure 3: P-values for comparison of EQCs and NEQCs at depth one and four from Figure 6.5 b) and d).

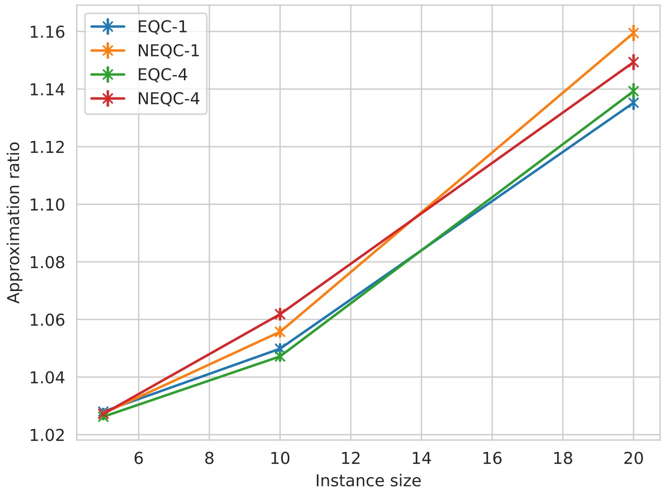

Figure 4: Mean and standard error of the mean for the one- and four-layer EQCs and NEQCs in Figure 6.5 b), d).

between the EQC at depths one and four is very small, and that increasing the circuit depth does not provide much benefit on this learning task.

Robustness of quantum reinforcement learning under hardware errors

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 $f : \mathbb{R}^M \to \mathbb{R}$ be the function under investigation, whose parameters $\boldsymbol{\theta} = (\theta_1, \ldots, \theta_M) \in \mathbb{R}^M$ are corrupted by a Gaussian noise $\theta_i \to \theta_i + \delta\theta_i$ with zero mean and variance σ^2, i.e.

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

Since the perturbations are independently distributed and Gaussian, all higher order moments can be evaluated starting from two-point correlators of the form $\mathbb{E}[\delta\theta_i\delta\theta_j]$, as dictated by *Wick's formulas* for multivariate normal distributions [305]

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

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 remain. For example, expression (2) for the fourth-order moment ($n = 4$) amounts to

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

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

$$f(\boldsymbol{\theta} + \delta\boldsymbol{\theta}) = f(\boldsymbol{\theta}) + \sum_{i=1}^{M} \frac{\partial f(\boldsymbol{\theta})}{\partial \theta_i} \delta\theta_i + \frac{1}{2!} \sum_{i,j=1}^{M} \frac{\partial^2 f(\boldsymbol{\theta})}{\partial \theta_i \partial \theta_j} \delta\theta_i \delta\theta_j + \frac{1}{3!} \sum_{i,j,k=1}^{M} \frac{\partial^3 f(\boldsymbol{\theta})}{\partial \theta_i \partial \theta_j \partial \theta_k} \delta\theta_i \delta\theta_j \delta\theta_k + \dots . \quad (4)$$

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 remainder 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. (1), (2) and (3) in (4), one can evaluate the expectation value of the function over the perturbations' distributions as

$$\begin{aligned} \mathbb{E}[f(\boldsymbol{\theta} + \delta\boldsymbol{\theta})] &= f(\boldsymbol{\theta}) + \partial_i f \, \mathbb{E}[\delta\theta_i] + \frac{1}{2} \partial_{ij} f \, \mathbb{E}[\delta\theta_i \delta\theta_j] + \frac{1}{3!} \partial_{ijk} f \, \mathbb{E}[\delta\theta_i \delta\theta_j \delta\theta_k] \\ &\quad + \frac{1}{4!} \partial_{ijkm} f \, \mathbb{E}[\delta\theta_i \delta\theta_j \delta\theta_k \delta\theta_m] + \dots \\ &= f(\boldsymbol{\theta}) + \frac{\sigma^2}{2} \partial_{ij} f \, \delta_{ij} + \frac{\sigma^4}{4!} \partial_{ijkm} f \, (\delta_{ij}\delta_{km} + \delta_{ik}\delta_{jm} + \delta_{im}\delta_{jk}) + \dots \\ &= f(\boldsymbol{\theta}) + \frac{\sigma^2}{2} \sum_i \frac{\partial^2 f}{\partial \theta_i^2} + \frac{\sigma^4}{4!} 3 \sum_{ij} \frac{\partial^4 f}{\partial \theta_i^2 \partial \theta_j^2} + \dots \end{aligned} \quad (5)$$

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

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

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

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

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

Thus, the full Taylor series can be formally written as

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

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

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

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

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 [**? ?**], one can show that any derivative of the function $f(\boldsymbol{\theta}) = \langle O \rangle = \mathrm{Tr}\big[O\, U(\boldsymbol{\theta})\, |0\rangle\!\langle 0|\, U^\dagger(\boldsymbol{\theta})\big]$ obeys

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

where $\|O\|_\infty$ is the infinity norm of the observable, namely its largest absolute eigenvalue. We give a proof of this below in Sec. 8.

Plugging this in Eq. (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

$$\begin{aligned}
\varepsilon_{\boldsymbol{\theta}} &= \left| \sum_{n=1}^{\infty} \frac{\sigma^{2n}}{(2n)!} (2n-1)!! \sum_{i_1,\ldots,i_n=1}^{M} \frac{\partial^{2n} f(\boldsymbol{\theta})}{\partial\theta_{i_1}^2 \ldots \partial\theta_{i_n}^2} \right| && (11) \\
&\leq \sum_{n=1}^{\infty} \frac{\sigma^{2n}}{(2n)!} (2n-1)!! \sum_{i_1,\ldots,i_n=1}^{M} \underbrace{\left| \frac{\partial^{2n} f(\boldsymbol{\theta})}{\partial\theta_{i_1}^2 \ldots \partial\theta_{i_n}^2} \right|}_{\leq \|O\|_\infty} \\
&\leq \sum_{n=1}^{\infty} \frac{\sigma^{2n}}{(2n)!} (2n-1)!! \, \|O\|_\infty \, M^n && (12) \\
&= \|O\|_\infty \sum_{n=1}^{\infty} \frac{1}{(2n)!} \frac{(2n)!}{2^n n!} (\sigma^2 M)^n = \|O\|_\infty \sum_{n=1}^{\infty} \frac{(M\sigma^2/2)^n}{n!} \\
&= \|O\|_\infty \left(e^{\sigma^2 M/2} - 1 \right) \\
&\implies \varepsilon_{\boldsymbol{\theta}} = |\mathbb{E}[f(\boldsymbol{\theta} + \delta\boldsymbol{\theta})] - f(\boldsymbol{\theta})| \leq \|O\|_\infty \left(e^{M\sigma^2/2} - 1 \right), && (13)
\end{aligned}$$

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 $f(\boldsymbol{\theta})$ and its perturbed version $f(\boldsymbol{\theta} + \delta\boldsymbol{\theta})$ cannot be larger than twice the maximum eigenvalue of O, $|f(\boldsymbol{\theta} + \delta\boldsymbol{\theta}) - f(\boldsymbol{\theta})| \leq |f(\boldsymbol{\theta} + \delta\boldsymbol{\theta})| + |f(\boldsymbol{\theta})| = 2\|O\|_\infty$, the bound (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. (8) are repeating the same calculation from above, one obtains

$$\left| \mathbb{E}[f(\boldsymbol{\theta} + \delta\boldsymbol{\theta})] - f(\boldsymbol{\theta}) - \frac{\sigma^2}{2} \operatorname{Tr}[H(\boldsymbol{\theta})] \right| \leq \|O\|_\infty \left(e^{M\sigma^2/2} - 1 - \frac{M\sigma^2}{2} \right). \qquad (14)$$

Parameter-Shift rule and bounds to the derivatives

Let $f(\boldsymbol{\theta}) = \operatorname{Tr}\big[O\, U(\boldsymbol{\theta}) \, |0\rangle\langle 0| \, U^\dagger(\boldsymbol{\theta})\big]$ be the expectation value of an observable O on the parametrized state $|\psi(\boldsymbol{\theta})\rangle = U(\boldsymbol{\theta})\,|0\rangle$ obtained with a parametrized quantum circuit $U(\boldsymbol{\theta})$. When the variational parameters $\boldsymbol{\theta} \in \mathbb{R}^M$ 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 [39, 33], as described in Section 2.2.1.1,

$$\frac{\partial f(\boldsymbol{\theta})}{\partial \theta_i} = \frac{1}{2}\Big(f\Big(\boldsymbol{\theta} + \frac{\pi}{2}\boldsymbol{e_i}\Big) - f\Big(\boldsymbol{\theta} - \frac{\pi}{2}\boldsymbol{e_i}\Big)\Big)\,, \tag{15}$$

where $\boldsymbol{e}_i$ 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 [280, 306]

$$\begin{aligned}
\frac{\partial^2 f(\boldsymbol{\theta})}{\partial \theta_i \theta_j} &= \frac{1}{2}\left[\frac{\partial}{\partial \theta_i} f\Big(\boldsymbol{\theta} + \frac{\pi}{2}\boldsymbol{e_j}\Big) - \frac{\partial}{\partial \theta_i} f\Big(\boldsymbol{\theta} - \frac{\pi}{2}\boldsymbol{e_j}\Big)\right] & (16)\\
&= \frac{1}{4}[f(\boldsymbol{\theta} + \frac{\pi}{2}\boldsymbol{e_j} + \frac{\pi}{2}\boldsymbol{e_i}) - f(\boldsymbol{\theta} + \frac{\pi}{2}\boldsymbol{e_j} - \frac{\pi}{2}\boldsymbol{e_i}) & (17)\\
&\quad - f\Big(\boldsymbol{\theta} - \frac{\pi}{2}\boldsymbol{e_j} + \frac{\pi}{2}\boldsymbol{e_i}\Big) + f\Big(\boldsymbol{\theta} - \frac{\pi}{2}\boldsymbol{e_j} - \frac{\pi}{2}\boldsymbol{e_i}\Big)]. & (18)
\end{aligned}$$

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

$$\begin{aligned}
\frac{\partial^2 f(\boldsymbol{\theta})}{\partial \theta_i^2} &= \frac{1}{4}[f(\boldsymbol{\theta} + \pi\boldsymbol{e_i}) - 2f(\boldsymbol{\theta}) + f(\boldsymbol{\theta} - \pi\boldsymbol{e_i})]\\
&= \frac{1}{2}[f(\boldsymbol{\theta} + \pi\boldsymbol{e_i}) - f(\boldsymbol{\theta})]\,,
\end{aligned} \tag{19}$$

where we used the fact that $f(\boldsymbol{\theta} + \pi\boldsymbol{e_i}) = f(\boldsymbol{\theta} - \pi\boldsymbol{e_i})$. 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 $\exp[-i\, 2\pi P/2] = \cos\pi\, \mathbb{I} - i \sin\pi\, P = -\mathbb{I}$, one has

$$\begin{aligned}
|\psi(\boldsymbol{\theta} - \pi\boldsymbol{e}_i)\rangle &= U_2 \exp[-i(\theta_i - \pi)P_i/2]\, U_1 \,|0\rangle\\
&= U_2 \exp[-i(\theta_i - \pi)P_i/2] \underbrace{-\exp[-i\, 2\pi\, P_i/2]}_{\mathbb{I}} U_1 \,|0\rangle\\
&= -U_2 \exp[-i(\theta_i - \pi + 2\pi)P_i/2] U_1 \,|0\rangle\\
&= -\,|\psi(\boldsymbol{\theta} + \pi\boldsymbol{e}_i)\rangle\;,
\end{aligned} \tag{20}$$

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

Hence, using Eq. (19) 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 combinations of circuit evaluations [280, 51]. Let $\boldsymbol{\alpha} = (\alpha_1, \ldots, \alpha_M) \in \mathbb{N}^M$ be a multi-index keeping track of the orders of derivatives, and let $|\boldsymbol{\alpha}| = \sum_{i=1}^{M} \alpha_i$. Then

$$\partial^{\boldsymbol{\alpha}} f(\boldsymbol{\theta}) := \frac{\partial^{|\boldsymbol{\alpha}|} f(\boldsymbol{\theta})}{\partial\theta_1^{\alpha_1} \ldots \partial\theta_M^{\alpha_M}} = \frac{1}{2^{|\boldsymbol{\alpha}|}} \sum_{m=1}^{2^{|\boldsymbol{\alpha}|}} s_m f(\tilde{\boldsymbol{\theta}}_m)\,, \tag{21}$$

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 $\|O\|_\infty = \max\{|o_i|\,, O = \sum_i o_i\, |o_i\rangle\!\langle o_i|\}$

$$|f(\boldsymbol{\theta})| = |\mathrm{Tr}[O\,\rho(\boldsymbol{\theta})]| \leq \|O\|_\infty \|\rho(\boldsymbol{\theta})\|_1 = \|O\|_\infty \quad \forall \boldsymbol{\theta} \in \mathbb{R}^M\,, \tag{22}$$

then one can bound the sum in Eq. (21) simply as

$$|\partial^{\boldsymbol{\alpha}} f(\boldsymbol{\theta})| \leq \frac{1}{2^{|\boldsymbol{\alpha}|}} \sum_{m=1}^{2^{|\boldsymbol{\alpha}|}} \left| f(\tilde{\boldsymbol{\theta}}_m) \right| \leq \|O\|_\infty\,. \tag{23}$$

Average value of the Hessian of random PQCs

In this section we derive the formulas (7.16) and (7.17) 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 $U(\boldsymbol{\theta}) \in \mathcal{U}(2^n)$, where $\mathcal{U}(2^n)$ is the group of unitary matrices of dimension 2^n. Given a set of parameter vectors $\{\boldsymbol{\theta}_1, \boldsymbol{\theta}_2, \ldots, \boldsymbol{\theta}_K\}$, one can construct the corresponding set of unitaries $\mathbb{U} = \{U_1, U_2, \ldots, U_K\}$, with $U_i = U(\boldsymbol{\theta}_i)$ 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) [48, 229, 44]. 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 [289, 307, 288, 48, 44]

$$\mathbb{E}_U[UAU^\dagger] = \int d\mu(U)\, UAU^\dagger = \frac{\mathbb{1}\operatorname{Tr}[A]}{2^n} \tag{24}$$

$$\mathbb{E}_U[AUBU^\dagger CUDU^\dagger] = \frac{\operatorname{Tr}[BD]\operatorname{Tr}[C]A + \operatorname{Tr}[B]\operatorname{Tr}[D]AC}{2^{2n}-1} \tag{25}$$

$$-\frac{\operatorname{Tr}[BD]AC + \operatorname{Tr}[B]\operatorname{Tr}[C]\operatorname{Tr}[D]A}{2^n(2^{2n}-1)} \tag{26}$$

Statistics of the Hessian

Let $f(\boldsymbol{\theta}) = \operatorname{Tr}\big[OU(\boldsymbol{\theta})\,|0\rangle\!\langle 0|\,U(\boldsymbol{\theta})^\dagger\big]$ and assume that the observable O is such that $\operatorname{Tr}[O] = 0$ and $\operatorname{Tr}\big[O^2\big] = 2^n$, as is the case of measuring a Pauli string. As shown in Eq. (19), diagonal elements of the Hessian matrix H can be calculated as

$$H_{ii} = \frac{\partial^2 f(\boldsymbol{\theta})}{\partial \theta_i^2} = \frac{1}{2}[f(\boldsymbol{\theta} + \pi \boldsymbol{e}_i) - f(\boldsymbol{\theta})]\,. \tag{27}$$

For simplicity, from now on we drop the explicit dependence on the parameter vector $\boldsymbol{\theta}$ 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 $U(\boldsymbol{\theta} + \pi\boldsymbol{e}_i)$ can be rewritten as

$$U(\boldsymbol{\theta} + \pi\boldsymbol{e}_i) = U_L e^{-i\pi P_i/2} U_R = -i\, U_L P_i U_R\,, \tag{28}$$

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

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

$$\mathbb{E}_{U_L}[f(\boldsymbol{\theta} + \pi\boldsymbol{e}_i)] = \mathbb{E}_{U_L}\Big[\operatorname{Tr}\Big[O\, U_L P_i U_R\,|0\rangle\!\langle 0|\, U_R^\dagger P_i U_L^\dagger\Big]\Big] \tag{29}$$

$$= \operatorname{Tr}\Big[O\,\mathbb{E}_{U_L}\Big[U_L P_i U_R\,|0\rangle\!\langle 0|\, U_R^\dagger P_i U_L^\dagger\Big]\Big] \tag{30}$$

$$= \operatorname{Tr}\left[O\,\frac{\operatorname{Tr}\big[P_i U_R\,|0\rangle\!\langle 0|\, U_R^\dagger P_i\big]\mathbb{1}}{2^n}\right] = \frac{\operatorname{Tr}[O]}{2^n} = 0\,, \tag{31}$$

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. (24) 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 $\mathbb{E}_{U_R}[f(\boldsymbol{\theta}+\pi\boldsymbol{e}_i)] = 0$. The same calculation for $f(\boldsymbol{\theta})$ shows that $\mathbb{E}_{U_R}[f(\boldsymbol{\theta})] = \mathbb{E}_{U_L}[f(\boldsymbol{\theta})] = 0$.

Thus, for every diagonal element of the Hessian, if either $\mathbb{U}_L$ or $\mathbb{U}_\mathbb{R}$ is a 1-design (that is Eq. (24) hold), then its expectation value vanishes

$$\mathbb{E}_{U_R,U_L}[H_{ii}] = 0 \quad \forall i \quad \text{if either } \mathbb{U}_L \text{ or } \mathbb{U}_\mathbb{R} \text{ is a 1-design.} \tag{32}$$

The variance of the diagonal elements can be calculated in a similar manner, even though the calculation is more involved. Substituting Eq. (27) in the defition of the variance, one obtains

$$\begin{aligned}\mathrm{Var}[H_{ii}] &:= \mathbb{E}[H_{ii}^2] - \mathbb{E}[H_{ii}]^2 = \mathbb{E}[H_{ii}^2] \\ &= \frac{1}{4}\left[\mathbb{E}[f(\boldsymbol{\theta}+\pi\boldsymbol{e}_i)^2] + \mathbb{E}[f(\boldsymbol{\theta})^2] - 2\mathbb{E}[f(\boldsymbol{\theta}+\pi\boldsymbol{e}_i)f(\boldsymbol{\theta})]\right].\end{aligned} \tag{33}$$

In order to use Eq. (26) for second moment integrals, we can rewrite these expectation values as follows

$$\begin{aligned}\mathbb{E}[f(\boldsymbol{\theta}+\pi\boldsymbol{e}_i)^2] &= \mathbb{E}\left[\mathrm{Tr}\left[O\, U_L P_i U_R \,|0\rangle\langle 0|\, U_R^\dagger P_i U_L^\dagger\right]^2\right] \\ &= \mathbb{E}\left[\mathrm{Tr}\left[O\, U_L P_i U_R \,|0\rangle\langle 0|\, U_R^\dagger P_i U_L^\dagger\right] \langle 0|U_R^\dagger P_i U_L^\dagger O U_L P_i U_R|0\rangle\right] \\ &= \mathbb{E}\left[\mathrm{Tr}\left[O\, U_L P_i U_R \,|0\rangle\langle 0|\, U_R^\dagger P_i U_L^\dagger O U_L P_i U_R \,|0\rangle\langle 0|\, U_R^\dagger P_i U_L^\dagger\right]\right] \\ &= \mathrm{Tr}\left[\mathbb{E}[O\, U_L P_i U_R \,|0\rangle\langle 0|\, U_R^\dagger P_i U_L^\dagger O U_L P_i U_R \,|0\rangle\langle 0|\, U_R^\dagger P_i U_L^\dagger]\right],\end{aligned} \tag{34}$$

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

$$\mathbb{E}_{U_L}[f(\boldsymbol{\theta}+\pi\boldsymbol{e}_i)^2] = \mathrm{Tr}\left[\mathbb{E}_{U_L}[O\, U_L \underbrace{P_i U_R \,|0\rangle\langle 0|\, U_R^\dagger P_i}_{B} U_L^\dagger O U_L \underbrace{P_i U_R \,|0\rangle\langle 0|\, U_R^\dagger P_i}_{B} U_L^\dagger]\right] \tag{35}$$

$$= \mathrm{Tr}\left[\frac{\mathrm{Tr}\left[B^2\right]\mathrm{Tr}[O]O + \mathrm{Tr}[B]^2 O^2}{2^{2n}-1} - \frac{\mathrm{Tr}\left[B^2\right]O^2 + \mathrm{Tr}[B]^2\,\mathrm{Tr}[O]O}{2^n(2^{2n}-1)}\right] \tag{36}$$

$$= \frac{\mathrm{Tr}[O]^2 + \mathrm{Tr}\left[O^2\right]}{2^{2n}-1} - \frac{\mathrm{Tr}\left[O^2\right] + \mathrm{Tr}[O]^2}{2^n(2^{2n}-1)} = \frac{1}{2^n+1}, \tag{37}$$

where in the second line we made use of Eq. (26), and the third line the used that $\mathrm{Tr}[B] = \mathrm{Tr}\left[B^2\right] = 1$ since $B = P_i U_R \,|0\rangle\langle 0|\, U_R^\dagger P_i$ is a projector, and that $\mathrm{Tr}[O] = 0$

and $\mathrm{Tr}\left[O^2\right] = 2^n$. Similarly, one can show that integration over $\mathbb{U}_R$ yields the same result. Also, the same calculation leads to $\mathbb{E}_{U_L}[f(\boldsymbol{\theta})^2] = \mathbb{E}_{U_R}[f(\boldsymbol{\theta})^2] = 1/(2^n+1)$. Thus, if either $\mathbb{U}_L$ or $\mathbb{U}_\mathbb{R}$ is a 2-design then

$$\mathbb{E}_{U_R,U_L}[f(\boldsymbol{\theta})^2] = \mathbb{E}_{U_R,U_L}[f(\boldsymbol{\theta}+\pi\boldsymbol{e}_i)^2] = \frac{1}{2^n+1} \quad \forall i \quad \text{if either } \mathbb{U}_L \text{ or } \mathbb{U}_\mathbb{R} \text{ is a 2-design.} \tag{38}$$

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

$$\begin{aligned}
\mathbb{E}_{U_L}[f(\boldsymbol{\theta}+\pi\boldsymbol{e}_i)f(\boldsymbol{\theta})] &= \mathrm{Tr}\left[\mathbb{E}_{U_L}\left[O\,U_L P_i U_R\,|0\rangle\!\langle 0|\,U_R^\dagger U_L^\dagger O U_L U_R\,|0\rangle\!\langle 0|\,U_R^\dagger P_i U_L^\dagger\right]\right] \\
&= \mathrm{Tr}\left[\frac{\mathrm{Tr}\left[P_i U_R\,|0\rangle\!\langle 0|\,U_R^\dagger\right]^2 O^2}{2^{2n}-1} - \frac{O^2}{2^n(2^{2n}-1)}\right] \\
&= \frac{1}{2^{2n}-1}\left[2^n\;\mathrm{Tr}\left[P_i U_R\,|0\rangle\!\langle 0|\,U_R^\dagger\right]^2 - 1\right].
\end{aligned} \tag{39}$$

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

$$\begin{aligned}
\mathbb{E}_{U_R}[f(\boldsymbol{\theta}+\pi\boldsymbol{e}_i)f(\boldsymbol{\theta})] &= \mathrm{Tr}\left[O\,U_L P_i\,\mathbb{E}_{U_R}\left[U_R\,|0\rangle\!\langle 0|\,U_R^\dagger U_L^\dagger O U_L U_R\,|0\rangle\!\langle 0|\,U_R^\dagger\right] P_i U_L^\dagger\right] \\
&= \mathrm{Tr}\left[O\,U_L P_i\,\frac{(2^n-1)U_L^\dagger O U_L}{2^n(2^{2n}-1)}\,P_i U_L^\dagger\right] \\
&= \frac{1}{2^n(2^n+1)}\,\mathrm{Tr}\left[O\,U_L P_i U_L^\dagger O U_L P_i U_L^\dagger\right].
\end{aligned} \tag{40}$$

If both of them are 2-designs, then continuing from Eq. (40), one obtains

$$\begin{aligned}
\mathbb{E}_{U_L,U_R}[f(\boldsymbol{\theta}+\pi\boldsymbol{e}_i)f(\boldsymbol{\theta})] &= \frac{1}{2^n(2^n+1)}\,\mathrm{Tr}\left[\mathbb{E}_{U_L}\left[O\,U_L P_i U_L^\dagger O U_L P_i U_L^\dagger\right]\right] \\
= \frac{1}{2^n(2^n+1)}\,\mathrm{Tr}&\left[\frac{\mathrm{Tr}[P_i]^2 O^2 + \mathrm{Tr}\left[P_i^2\right]\mathrm{Tr}[O]O}{2^{2n}-1} - \frac{\mathrm{Tr}\left[P_i^2\right]O^2 + \mathrm{Tr}[P_i]^2\,\mathrm{Tr}[O]O}{2^n(2^{2n}-1)}\right] \\
= -\frac{1}{2^n(2^n+1)}\,\frac{\mathrm{Tr}\left[P_i^2\right]\mathrm{Tr}\left[O^2\right]}{2^n(2^{2n}-1)} &= -\frac{1}{(2^n+1)(2^{2n}-1)} \in \mathcal{O}\left(2^{-3n}\right)
\end{aligned} \tag{41}$$

Finally, plugging Eqs. (39), (40) and (41) in Eq. (33), one has $\forall i = 1, \dots, M$

$$
\begin{aligned}
\mathrm{Var}_{U_L,U_R}[H_{ii}] &= \frac{1}{2}\mathbb{E}[f(\boldsymbol{\theta})^2] - \frac{1}{2}\mathbb{E}[f(\boldsymbol{\theta}+\pi\boldsymbol{e}_i)f(\boldsymbol{\theta})] \\
&= \frac{1}{2(2^n+1)} - \frac{1}{2}\begin{cases} \frac{1}{2^{2n}-1}\left[2^n \ \mathrm{Tr}\left[P_i U_R \left|0\right\rangle\!\left\langle 0\right| U_R^\dagger\right]^2 - 1\right] & \forall i\,, \text{if } \mathbb{U}_L \text{ 2-design} \\ \frac{1}{2^n(2^n+1)}\,\mathrm{Tr}\left[O\,U_L P_i U_L^\dagger O U_L P_i U_L^\dagger\right] & \forall i\,, \text{if } \mathbb{U}_R \text{ 2-design} \\ -\frac{1}{(2^n+1)(2^{2n}-1)} & \forall i\,, \text{if } \mathbb{U}_L, \mathbb{U}_R \text{ 2-designs} \end{cases}
\end{aligned} \tag{42}
$$

where $\mathbb{U}_R = \mathbb{U}_R^{(i)}$ and $\mathbb{U}_L = \mathbb{U}_L^{(i)}$ are defined as in Eq. (28) 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 [51, 48], as from Eq. (42) one can check that $\mathrm{Var}[H_{ii}] \in \mathcal{O}(2^{-n})$.

Statistics of the trace of the Hessian

The average value of the trace of the Hessian is easily found to be zero using Eq. (32), in fact

$$
\mathbb{E}_{U_R,U_L}[\mathrm{Tr}[H]] = \sum_{i=1}^{M} \mathbb{E}_{U_R^{(i)},U_L^{(i)}}[H_{ii}] = 0\,, \tag{43}
$$

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

$$
\mathrm{Var}_{U_R,U_L}[\mathrm{Tr}[H]] = \mathrm{Var}\left[\sum_{i=1}^{M} H_{ii}\right] = \sum_{i=1}^{M}\mathrm{Var}[H_{ii}] + 2\sum_{i<j}^{M}\mathrm{Cov}[H_{ii}H_{jj}]\,. \tag{44}
$$

We can upper bound this quantity using the covariance inequality [308],

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

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

$$
\mathrm{Var}_{U_R,U_L}[\mathrm{Tr}[H]] \leq \sum_{i=1}^{M}\mathrm{Var}[H_{ii}] + 2\sum_{i<j}^{M}\mathrm{Var}[H_{ii}] \in \mathcal{O}\left(\frac{M^2}{2^n}\right). \tag{45}
$$

Alternatively, one can obtain a tighter yet qualitative approximation by explicitly considering the nature of the sums in Eq. (44). First, by using Eq. (27), the covariance term is explicitly

$$\begin{aligned}\mathrm{Cov}[H_{ii}, H_{jj}] &= \mathbb{E}[H_{ii}H_{jj}] \\ &= \frac{1}{4}\mathbb{E}[(f_i - f)(f_j - f)] \\ &= \frac{1}{4}\mathbb{E}[f^2] + \frac{1}{4}\mathbb{E}[f_i f_j] - \frac{1}{4}\mathbb{E}[f_i f] - \frac{1}{4}\mathbb{E}[f_j f]\,,\end{aligned} \tag{46}$$

where for ease of notation we defined $f_{i,j} = f(\boldsymbol{\theta} + \pi \boldsymbol{e}_{i,j})$ and $f = f(\boldsymbol{\theta})$. 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

$$|\mathbb{E}[f_i f_j]| \leq \sqrt{\mathbb{E}[f_i^2]\mathbb{E}[f_j^2]} = \frac{1}{2^n + 1} \quad \text{and} \quad |\mathbb{E}[f_i f]| \leq \sqrt{\mathbb{E}[f_i^2]\mathbb{E}[f^2]} = \frac{1}{2^n + 1}\,, \tag{47}$$

where we have used $E[f^2] = E[f_i^2] = 1/(2^n + 1)$ from Eq. (38). Then, the variance can be written as

$$\begin{aligned}&\mathrm{Var}_{U_R,U_L}[\mathrm{Tr}[H]] = \sum_{i=1}^{M} \mathrm{Var}[H_{ii}] + 2\sum_{i<j}^{M} \mathbb{E}[H_{ii}H_{jj}] \\ &= \sum_{i=1}^{M} \frac{\mathbb{E}[f^2] - \mathbb{E}[f_i f]}{2} + 2\sum_{i<j}^{M} \frac{\mathbb{E}[f^2] + \mathbb{E}[f_i f_j] - \mathbb{E}[f_i f] - \mathbb{E}[f_j f]}{4} \\ &= \frac{1}{2}\left(\sum_{i=1}^{M} + \sum_{i<j}^{M}\right)\mathbb{E}[f^2] - \frac{1}{2}\left(\sum_{i=1}^{M}\mathbb{E}[f_i f] + \sum_{i<j}^{M}\mathbb{E}[f_i f] + \sum_{i<j}^{M}\mathbb{E}[f_j f]\right) + \frac{1}{2}\sum_{i<j}^{M}\mathbb{E}[f_i f_j] \\ &= \frac{M(M+1)}{4}\mathbb{E}[f^2] \underbrace{- \frac{M}{2}\sum_{i=1}^{M}\mathbb{E}[f_i f] + \frac{1}{2}\sum_{i<j}^{M}\mathbb{E}[f_i f_j]}_{\Delta}\,.\end{aligned} \tag{48}$$

Numerical simulations In addition to Figure 7.6 in the main text, in Figure 5 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 $\mathrm{Tr}[H(\boldsymbol{\theta})]$ 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

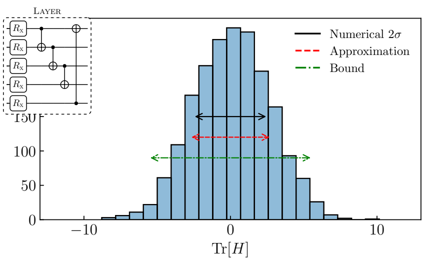

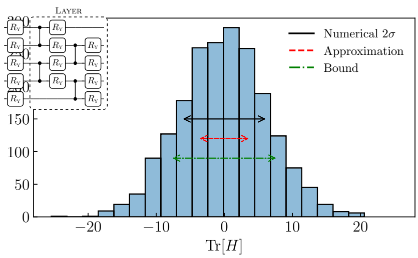

Figure 5: Simulation results of evaluating the trace of the Hessian matrix for two different hardware-efficient ansatzes with random values of the parameters. The plot on the left is obtained using the layer template shown in the figure for $n = 6$ qubits and $l = 6$ layers. The plot on the right instead with $n = 5$ and $l = 5$ layers of the template shown in the corresponding inset. The simulations are performed by sampling 2000 random parameter vectors $\boldsymbol{\theta}_m$ with $\theta_i \sim \mathrm{Unif}[0, 2\pi[$, evaluating the trace of the Hessian matrix $\mathrm{Tr}[H(\boldsymbol{\theta})]$, and then building the histogram to show its frequency distribution. In both experiments the measured observable is $Z^{\otimes n}$. 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. (49), "Bound" (dashed-dotted green) twice the square root of the upper Bound in Eq. (45). These parametrized circuits correspond to the templates `BasicEntanglinLayer` and `Simplified2Design` defined in Pennylane [**?**], and used for example in [44] to study barren plateaus.

the Eq. (48) with $\Delta = 0$, “Bound" (dashed-dotted green) twice the square root of the upper Bound in Eq. (45).

All simulations confirm the bound (45), and, more interestingly, both the circuit on the left of Fig. 5 and the one in Fig. 7.6 in the main text, have a numerical variance which is very well approximated by Eq. (48) with $\Delta = 0$. We conjecture this is due to the fact that all correlation terms in Eq. (48) are roughly of the same order of magnitude (see Eq. (47)), 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

$$\mathrm{Var}_{U_R,U_L}[\mathrm{Tr}[H]] \approx \frac{M(M+1)}{4}\mathbb{E}[f^2] = \frac{M(M+1)}{4(2^n+1)} \approx \frac{1}{4}\frac{M^2}{2^n}\,, \tag{49}$$

which is four times smaller then the upper bound Eq. (45), 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.

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

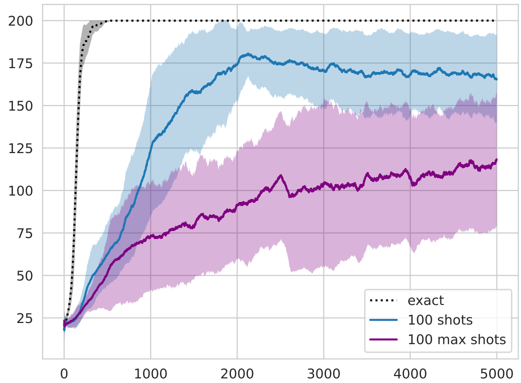

Figure 6: Performance of agents trained with a fixed number of 100 shots (blue) and $m_{\max} = 100$ with flexible shot allocation (purple), compared to model trained without shot noise (black dotted curve).

Visualization of CartPole policies obtained with Q-learning

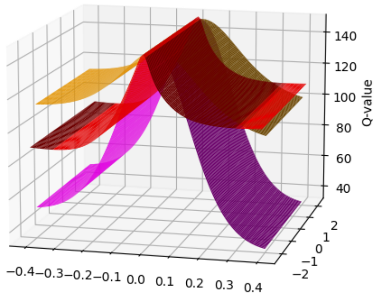

(a) $\sigma = 0$

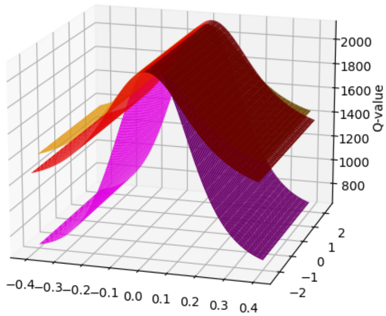

(b) $\sigma = 0.2$

Figure 7: Visualization of the Q-functions learned in the noise-free (a) and noisy (b) settings. The red surface shows Q-values for pole angle and cart position, orange for pole angle and cart velocity, and magenta for pole angle and pole velocity.

Bibliography

[1] Peter W Shor. Polynomial-time algorithms for prime factorization and discrete logarithms on a quantum computer. *SIAM review*, 41(2):303–332, 1999.

[2] Frank Arute, Kunal Arya, Ryan Babbush, Dave Bacon, Joseph C Bardin, Rami Barends, Rupak Biswas, Sergio Boixo, Fernando GSL Brandao, David A Buell, et al. Quantum supremacy using a programmable superconducting processor. *Nature*, 574(7779):505–510, 2019.

[3] Han-Sen Zhong, Hui Wang, Yu-Hao Deng, Ming-Cheng Chen, Li-Chao Peng, Yi-Han Luo, Jian Qin, Dian Wu, Xing Ding, Yi Hu, et al. Quantum computational advantage using photons. *Science*, 370(6523):1460–1463, 2020.

[4] Lov K Grover. A fast quantum mechanical algorithm for database search. In *Proceedings of the twenty-eighth annual ACM symposium on Theory of computing*, pages 212–219, 1996.

[5] Aram W Harrow, Avinatan Hassidim, and Seth Lloyd. Quantum algorithm for linear systems of equations. *Physical review letters*, 103(15):150502, 2009.

[6] Seth Lloyd, Masoud Mohseni, and Patrick Rebentrost. Quantum algorithms for supervised and unsupervised machine learning. *arXiv preprint arXiv:1307.0411*, 2013.

[7] Patrick Rebentrost, Masoud Mohseni, and Seth Lloyd. Quantum support vector machine for big data classification. *Physical review letters*, 113(13):130503, 2014.

[8] Vittorio Giovannetti, Seth Lloyd, and Lorenzo Maccone. Quantum random access memory. *Physical review letters*, 100(16):160501, 2008.

[9] Scott Aaronson. Read the fine print. *Nature Physics*, 11(4):291–293, 2015.

[10] Ewin Tang. A quantum-inspired classical algorithm for recommendation systems. In *Proceedings of the 51st annual ACM SIGACT symposium on theory of computing*, pages 217–228, 2019.

[11] Ainesh Bakshi and Ewin Tang. An improved classical singular value transformation for quantum machine learning. *arXiv preprint arXiv:2303.01492*, 2023.

[12] John Preskill. Quantum computing in the nisq era and beyond. *Quantum*, 2:79, 2018.

[13] Kishor Bharti, Alba Cervera-Lierta, Thi Ha Kyaw, Tobias Haug, Sumner Alperin-Lea, Abhinav Anand, Matthias Degroote, Hermanni Heimonen, Jakob S Kottmann, Tim Menke, et al. Noisy intermediate-scale quantum (nisq) algorithms. *arXiv preprint arXiv:2101.08448*, 2021.

[14] Marco Cerezo, Andrew Arrasmith, Ryan Babbush, Simon C Benjamin, Suguru Endo, Keisuke Fujii, Jarrod R McClean, Kosuke Mitarai, Xiao Yuan, Lukasz Cincio, et al. Variational quantum algorithms. *Nature Reviews Physics*, 3(9):625–644, 2021.

[15] E. Farhi, J. Goldstone, and S. Gutmann. A Quantum Approximate Optimization Algorithm. *Preprint at arXiv:1411.4028*, 2014.

[16] Zhihui Wang, Stuart Hadfield, Zhang Jiang, and Eleanor G Rieffel. Quantum approximate optimization algorithm for maxcut: A fermionic view. *Physical Review A*, 97(2):022304, 2018.

[17] Stuart Hadfield, Zhihui Wang, Bryan O'Gorman, Eleanor Rieffel, Davide Venturelli, Rupak Biswas, Stuart Hadfield, Zhihui Wang, Bryan O'Gorman, Eleanor G. Rieffel, Davide Venturelli, and Rupak Biswas. From the Quantum Approximate Optimization Algorithm to a Quantum Alternating Operator Ansatz. *Algorithms*, 12(2):34, feb 2019.

[18] Alberto Peruzzo, Jarrod McClean, Peter Shadbolt, Man-Hong Yung, Xiao-Qi Zhou, Peter J Love, Alán Aspuru-Guzik, and Jeremy L O'brien. A variational eigenvalue solver on a photonic quantum processor. *Nature communications*, 5:4213, 2014.

[19] M-H Yung, J Casanova, Antonio Mezzacapo, J McClean, L Lamata, A Aspuru-Guzik, and E Solano. From transistor to trapped-ion computers for quantum chemistry. *Scientific Reports*, 4(3589):1–7, 2014.

[20] Jarrod R McClean, Jonathan Romero, Ryan Babbush, and Alán Aspuru-Guzik. The theory of variational hybrid quantum-classical algorithms. *New Journal of Physics*, 18(2):023023, 2016.

[21] P. J. J. O'Malley, R. Babbush, I. D. Kivlichan, J. Romero, J. R. McClean, R. Barends, J. Kelly, P. Roushan, A. Tranter, N. Ding, et al. Scalable quantum simulation of molecular energies. *Physical Review X*, 6(3):031007, 2016.

[22] Raffaele Santagati, Jianwei Wang, Antonio A. Gentile, Stefano Paesani, Nathan Wiebe, Jarrod R. McClean, Sam Morley-Short, Peter J. Shadbolt, Damien Bonneau, Joshua W. Silverstone, David P. Tew, Xiaoqi Zhou, Jeremy L. O'Brien, and Mark G. Thompson. Witnessing eigenstates for quantum simulation of hamiltonian spectra. *Science Advances*, 4:1, 2018.

[23] Jonathan Romero, Jonathan P Olson, and Alan Aspuru-Guzik. Quantum autoencoders for efficient compression of quantum data. *Quantum Science and Technology*, 2(4):045001, 2017.

[24] Edward Grant, Marcello Benedetti, Shuxiang Cao, Andrew Hallam, Joshua Lockhart, Vid Stojevic, Andrew G. Green, and Simone Severini. Hierarchical quantum classifiers. *npj Quantum Information*, 4(1):65, dec 2018.

[25] Jin-Guo Liu and Lei Wang. Differentiable learning of quantum circuit Born machines. *Physical Review A*, 98(6):062324, dec 2018.

[26] Vojtěch Havlíček, Antonio D Córcoles, Kristan Temme, Aram W Harrow, Abhinav Kandala, Jerry M Chow, and Jay M Gambetta. Supervised learning with quantum-enhanced feature spaces. *Nature*, 567(7747):209–212, 2019.

[27] Marcello Benedetti, Edward Grant, Leonard Wossnig, and Simone Severini. Adversarial quantum circuit learning for pure state approximation. *New Journal of Physics*, 21(4):043023, apr 2019.

[28] Brian Coyle, Daniel Mills, Vincent Danos, and Elham Kashefi. The born supremacy: Quantum advantage and training of an ising born machine. *npj Quantum Information*, 6(1):1–11, 2020.

[29] David Silver, Aja Huang, Chris J Maddison, Arthur Guez, Laurent Sifre, George Van Den Driessche, Julian Schrittwieser, Ioannis Antonoglou, Veda Panneershelvam, Marc Lanctot, et al. Mastering the game of go with deep neural networks and tree search. *nature*, 529(7587):484–489, 2016.

[30] John Jumper, Richard Evans, Alexander Pritzel, Tim Green, Michael Figurnov, Olaf Ronneberger, Kathryn Tunyasuvunakool, Russ Bates, Augustin Žídek, Anna Potapenko, et al. Highly accurate protein structure prediction with alphafold. *Nature*, 596(7873):583–589, 2021.

[31] Kathryn Tunyasuvunakool, Jonas Adler, Zachary Wu, Tim Green, Michal Zielinski, Augustin Žídek, Alex Bridgland, Andrew Cowie, Clemens Meyer, Agata Laydon, et al. Highly accurate protein structure prediction for the human proteome. *Nature*, 596(7873):590–596, 2021.

[32] Aditya Ramesh, Prafulla Dhariwal, Alex Nichol, Casey Chu, and Mark Chen. Hierarchical text-conditional image generation with clip latents. *arXiv preprint arXiv:2204.06125*, 2022.

[33] Maria Schuld, Ville Bergholm, Christian Gogolin, Josh Izaac, and Nathan Killoran. Evaluating analytic gradients on quantum hardware. *Physical Review A*, 99(3):032331, mar 2019.

[34] Kevin J Sung, Jiahao Yao, Matthew P Harrigan, Nicholas C Rubin, Zhang Jiang, Lin Lin, Ryan Babbush, and Jarrod R McClean. Using models to improve optimizers for variational quantum algorithms. *Quantum Science and Technology*, 5(4):044008, 2020.

[35] Xavier Bonet-Monroig, Hao Wang, Diederick Vermetten, Bruno Senjean, Charles Moussa, Thomas Bäck, Vedran Dunjko, and Thomas E O'Brien. Performance comparison of optimization methods on variational quantum algorithms. *arXiv preprint arXiv:2111.13454*, 2021.

[36] Aidan Pellow-Jarman, Ilya Sinayskiy, Anban Pillay, and Francesco Petruccione. A comparison of various classical optimizers for a variational quantum linear solver. *Quantum Information Processing*, 20(6):1–14, 2021.

[37] Samuel Yen-Chi Chen, Chih-Min Huang, Chia-Wei Hsing, Hsi-Sheng Goan, and Ying-Jer Kao. Variational quantum reinforcement learning via evolutionary optimization. *Machine Learning: Science and Technology*, 3(1):015025, 2022.

[38] Diederik P. Kingma and Jimmy Ba. Adam: A Method for Stochastic Optimization. In *International Conference on Learning Representations*, 2015.

[39] Kosuke Mitarai, Makoto Negoro, Masahiro Kitagawa, and Keisuke Fujii. Quantum circuit learning. *Physical Review A*, 98(3):032309, 2018.

[40] Gavin E Crooks. Gradients of parameterized quantum gates using the parameter-shift rule and gate decomposition. *arXiv preprint arXiv:1905.13311*, 2019.

[41] Jarrod R. McClean, Sergio Boixo, Vadim N. Smelyanskiy, Ryan Babbush, and Hartmut Neven. Barren plateaus in quantum neural network training landscapes. *Nature Communications*, 9(1):4812, dec 2018.

[42] Aram W. Harrow and Richard A. Low. Random Quantum Circuits are Approximate 2-designs. *Communications in Mathematical Physics*, 291(1):257–302, oct 2009.

[43] Sergio Boixo, Sergei V. Isakov, Vadim N. Smelyanskiy, Ryan Babbush, Nan Ding, Zhang Jiang, Michael J. Bremner, John M. Martinis, and Hartmut Neven. Characterizing quantum supremacy in near-term devices. *Nature Physics*, 14(6):595–600, jun 2018.

[44] Marco Cerezo, Akira Sone, Tyler Volkoff, Lukasz Cincio, and Patrick J Coles. Cost function dependent barren plateaus in shallow parametrized quantum circuits. *Nature communications*, 12(1):1–12, 2021.

[45] AV Uvarov and Jacob D Biamonte. On barren plateaus and cost function locality in variational quantum algorithms. *Journal of Physics A: Mathematical and Theoretical*, 54(24):245301, 2021.

[46] Zoë Holmes, Andrew Arrasmith, Bin Yan, Patrick J Coles, Andreas Albrecht, and Andrew T Sornborger. Barren plateaus preclude learning scramblers, arxiv preprints. *arXiv preprint arXiv:2009.14808*, 2020.

[47] Carlos Ortiz Marrero, Mária Kieferová, and Nathan Wiebe. Entanglement-induced barren plateaus. *PRX Quantum*, 2(4):040316, 2021.

[48] Zoë Holmes, Kunal Sharma, Marco Cerezo, and Patrick J Coles. Connecting ansatz expressibility to gradient magnitudes and barren plateaus. *PRX Quantum*, 3(1):010313, 2022.

[49] Samson Wang, Enrico Fontana, Marco Cerezo, Kunal Sharma, Akira Sone, Lukasz Cincio, and Patrick J Coles. Noise-induced barren plateaus in variational quantum algorithms. *Nature communications*, 12(1):1–11, 2021.

[50] Daniel Stilck França and Raul Garcia-Patron. Limitations of optimization algorithms on noisy quantum devices. *Nature Physics*, 17(11):1221–1227, 2021.

[51] Marco Cerezo and Patrick J Coles. Higher order derivatives of quantum neural networks with barren plateaus. *Quantum Science and Technology*, 6(3):035006, 2021.

[52] Andrew Arrasmith, M Cerezo, Piotr Czarnik, Lukasz Cincio, and Patrick J Coles. Effect of barren plateaus on gradient-free optimization. *Quantum*, 5:558, 2021.

[53] Edward Grant, Leonard Wossnig, Mateusz Ostaszewski, and Marcello Benedetti. An initialization strategy for addressing barren plateaus in parametrized quantum circuits. *Quantum*, 3:214, 2019.

[54] Tyler Volkoff and Patrick J Coles. Large gradients via correlation in random parameterized quantum circuits. *Quantum Science and Technology*, 6(2):025008, 2021.

[55] Iris Cong, Soonwon Choi, and Mikhail D Lukin. Quantum convolutional neural networks. *Nature Physics*, 15(12):1273–1278, 2019.

[56] Arthur Pesah, M Cerezo, Samson Wang, Tyler Volkoff, Andrew T Sornborger, and Patrick J Coles. Absence of barren plateaus in quantum convolutional neural networks. *Physical Review X*, 11(4):041011, 2021.

[57] Louis Schatzki, Martin Larocca, Frederic Sauvage, and M Cerezo. Theoretical guarantees for permutation-equivariant quantum neural networks. *arXiv preprint arXiv:2210.09974*, 2022.

[58] Stefan H Sack, Raimel A Medina, Alexios A Michailidis, Richard Kueng, and Maksym Serbyn. Avoiding barren plateaus using classical shadows. *PRX Quantum*, 3(2):020365, 2022.

[59] Edward Farhi, Jeffrey Goldstone, and Sam Gutmann. A quantum approximate optimization algorithm. *arXiv preprint arXiv:1411.4028*, 2014.

[60] Gavin E Crooks. Performance of the quantum approximate optimization algorithm on the maximum cut problem. *arXiv preprint arXiv:1811.08419*, 2018.

[61] Fernando G. Brandao, Michael Broughton, Edward Farhi, Sam Gutmann, and Hartmut Neven. For fixed control parameters the quantum approximate optimization algorithm's objective function value concentrates for typical instances. *arXiv preprint arXiv:1812.04170*, 2018.

[62] Stuart Hadfield, Zhihui Wang, Bryan O'Gorman, Eleanor G Rieffel, Davide Venturelli, and Rupak Biswas. From the quantum approximate optimization algorithm to a quantum alternating operator ansatz. *Algorithms*, 12(2):34, 2019.

[63] Gian Giacomo Guerreschi and Anne Y Matsuura. Qaoa for max-cut requires hundreds of qubits for quantum speed-up. *Scientific reports*, 9(1):1–7, 2019.

[64] Román Orús, Samuel Mugel, and Enrique Lizaso. Quantum computing for finance: Overview and prospects. *Reviews in Physics*, 4:100028, 2019.

[65] Michael Streif, Sheir Yarkoni, Andrea Skolik, Florian Neukart, and Martin Leib. Beating classical heuristics for the binary paint shop problem with the quantum approximate optimization algorithm. *Physical Review A*, 104(1):012403, 2021.

[66] Andre Luckow, Johannes Klepsch, and Josef Pichlmeier. Quantum computing: Towards industry reference problems. *Digitale Welt*, 5(2):38–45, 2021.

[67] Constantin Dalyac, Loïc Henriet, Emmanuel Jeandel, Wolfgang Lechner, Simon Perdrix, Marc Porcheron, and Margarita Veshchezerova. Qualifying quantum approaches for hard industrial optimization problems. a case study in the field of smart-charging of electric vehicles. *EPJ Quantum Technology*, 8(1):12, 2021.

[68] Matthew P Harrigan, Kevin J Sung, Matthew Neeley, Kevin J Satzinger, Frank Arute, Kunal Arya, Juan Atalaya, Joseph C Bardin, Rami Barends, Sergio Boixo, et al. Quantum approximate optimization of non-planar graph problems on a planar superconducting processor. *Nature Physics*, 17(3):332–336, 2021.

[69] Sepehr Ebadi, Alexander Keesling, Madelyn Cain, Tout T Wang, Harry Levine, Dolev Bluvstein, Giulia Semeghini, Ahmed Omran, J-G Liu, Rhine Samajdar, et al. Quantum optimization of maximum independent set using rydberg atom arrays. *Science*, page eabo6587, 2022.

[70] Ying Li and Simon C Benjamin. Efficient variational quantum simulator incorporating active error minimization. *Physical Review X*, 7(2):021050, 2017.

[71] Harper R Grimsley, Sophia E Economou, Edwin Barnes, and Nicholas J Mayhall. An adaptive variational algorithm for exact molecular simulations on a quantum computer. *Nature communications*, 10(1):1–9, 2019.

[72] Arthur G Rattew, Shaohan Hu, Marco Pistoia, Richard Chen, and Steve Wood. A domain-agnostic, noise-resistant, hardware-efficient evolutionary variational quantum eigensolver. *arXiv preprint arXiv:1910.09694*, 2019.

[73] Google AI Quantum, Collaborators*†, Frank Arute, Kunal Arya, Ryan Babbush, Dave Bacon, Joseph C Bardin, Rami Barends, Sergio Boixo, Michael Broughton, Bob B Buckley, et al. Hartree-fock on a superconducting qubit quantum computer. *Science*, 369(6507):1084–1089, 2020.

[74] Maria Schuld, Alex Bocharov, Krysta M Svore, and Nathan Wiebe. Circuit-centric quantum classifiers. *Physical Review A*, 101(3):032308, 2020.

[75] Andrea Skolik, Sofiene Jerbi, and Vedran Dunjko. Quantum agents in the gym: a variational quantum algorithm for deep q-learning. *Quantum*, 6:720, 2022.

[76] André Sequeira, Luis Paulo Santos, and Luís Soares Barbosa. Variational quantum policy gradients with an application to quantum control. *arXiv preprint arXiv:2203.10591*, 2022.

[77] Asel Sagingalieva, Andrii Kurkin, Artem Melnikov, Daniil Kuhmistrov, Michael Perelshtein, Alexey Melnikov, Andrea Skolik, and David Von Dollen. Hyperparameter optimization of hybrid quantum neural networks for car classification. *arXiv preprint arXiv:2205.04878*, 2022.

[78] Kristan Temme, Sergey Bravyi, and Jay M Gambetta. Error mitigation for short-depth quantum circuits. *Physical review letters*, 119(18):180509, 2017.

[79] Suguru Endo, Simon C Benjamin, and Ying Li. Practical quantum error mitigation for near-future applications. *Physical Review X*, 8(3):031027, 2018.

[80] Philippe Suchsland, Francesco Tacchino, Mark H Fischer, Titus Neupert, Panagiotis Kl Barkoutsos, and Ivano Tavernelli. Algorithmic error mitigation scheme for current quantum processors. *Quantum*, 5:492, 2021.

[81] Ryuji Takagi, Suguru Endo, Shintaro Minagawa, and Mile Gu. Fundamental limits of quantum error mitigation. *npj Quantum Information*, 8(1):1–11, 2022.

[82] Yihui Quek, Daniel Stilck França, Sumeet Khatri, Johannes Jakob Meyer, and Jens Eisert. Exponentially tighter bounds on limitations of quantum error mitigation. *arXiv preprint arXiv:2210.11505*, 2022.

[83] Christopher M Bishop and Nasser M Nasrabadi. *Pattern recognition and machine learning*, volume 4. Springer, 2006.

[84] Kevin P. Murphy. *Probabilistic Machine Learning: An introduction*. MIT Press, 2022.

[85] Ian Goodfellow, Yoshua Bengio, and Aaron Courville. *Deep Learning*. MIT Press, 2016. http://www.deeplearningbook.org.

[86] Henry J Kelley. Gradient theory of optimal flight paths. *Ars Journal*, 30(10):947–954, 1960.

[87] David E. Rumelhart, Geoffrey E. Hinton, and Ronald J. Williams. Learning representations by back-propagating errors. *Nature*, 323(6088):533–536, oct 1986.

[88] Sepp Hochreiter. The vanishing gradient problem during learning recurrent neural nets and problem solutions. *International Journal of Uncertainty, Fuzziness and Knowledge-Based Systems*, 6(02):107–116, 1998.

[89] Siddharth Krishna Kumar. On weight initialization in deep neural networks. *arXiv preprint arXiv:1704.08863*, 2017.

[90] Xavier Glorot, Antoine Bordes, and Yoshua Bengio. Deep sparse rectifier neural networks. In *Proceedings of the fourteenth international conference on artificial intelligence and statistics*, pages 315–323. JMLR Workshop and Conference Proceedings, 2011.

[91] Alex Graves and Jürgen Schmidhuber. Offline handwriting recognition with multidimensional recurrent neural networks. *Advances in neural information processing systems*, 21, 2008.

[92] George Cybenko. Approximation by superpositions of a sigmoidal function. *Mathematics of control, signals and systems*, 2(4):303–314, 1989.

[93] Kurt Hornik. Approximation capabilities of multilayer feedforward networks. *Neural networks*, 4(2):251–257, 1991.

[94] Ken-Ichi Funahashi. On the approximate realization of continuous mappings by neural networks. *Neural networks*, 2(3):183–192, 1989.

[95] Kurt Hornik, Maxwell Stinchcombe, and Halbert White. Multilayer feedforward networks are universal approximators. *Neural networks*, 2(5):359–366, 1989.

[96] Zhou Lu, Hongming Pu, Feicheng Wang, Zhiqiang Hu, and Liwei Wang. The expressive power of neural networks: A view from the width. *Advances in neural information processing systems*, 30, 2017.

[97] Vitaly Maiorov and Allan Pinkus. Lower bounds for approximation by mlp neural networks. *Neurocomputing*, 25(1-3):81–91, 1999.

[98] Stuart Geman, Elie Bienenstock, and René Doursat. Neural networks and the bias/variance dilemma. *Neural computation*, 4(1):1–58, 1992.

[99] Richard Bellman. Dynamic programming. *Science*, 153(3731):34–37, 1966.

[100] Michael M Bronstein, Joan Bruna, Taco Cohen, and Petar Veličković. Geometric deep learning: Grids, groups, graphs, geodesics, and gauges. *arXiv preprint arXiv:2104.13478*, 2021.

[101] David J Gross. The role of symmetry in fundamental physics. *Proceedings of the National Academy of Sciences*, 93(25):14256–14259, 1996.

[102] Kunihiko Fukushima and Sei Miyake. Neocognitron: A self-organizing neural network model for a mechanism of visual pattern recognition. In *Competition and cooperation in neural nets*, pages 267–285. Springer, 1982.

[103] Yann LeCun, Bernhard Boser, John S Denker, Donnie Henderson, Richard E Howard, Wayne Hubbard, and Lawrence D Jackel. Backpropagation applied to handwritten zip code recognition. *Neural computation*, 1(4):541–551, 1989.

[104] Yann LeCun, Yoshua Bengio, et al. Convolutional networks for images, speech, and time series. *The handbook of brain theory and neural networks*, 3361(10):1995, 1995.

[105] Franco Scarselli, Marco Gori, Ah Chung Tsoi, Markus Hagenbuchner, and Gabriele Monfardini. The graph neural network model. *IEEE transactions on neural networks*, 20(1):61–80, 2008.

[106] Elman Mansimov, Omar Mahmood, Seokho Kang, and Kyunghyun Cho. Molecular geometry prediction using a deep generative graph neural network. *Scientific reports*, 9(1):1–13, 2019.

[107] Zonghan Wu, Shirui Pan, Fengwen Chen, Guodong Long, Chengqi Zhang, and S Yu Philip. A comprehensive survey on graph neural networks. *IEEE transactions on neural networks and learning systems*, 32(1):4–24, 2020.

[108] Quentin Cappart, Didier Chételat, Elias Khalil, Andrea Lodi, Christopher Morris, and Petar Veličković. Combinatorial optimization and reasoning with graph neural networks. *arXiv preprint arXiv:2102.09544*, 2021.

[109] Richard S Sutton and Andrew G Barto. *Reinforcement learning: An introduction*. MIT press, 2018.

[110] Francisco S Melo. Convergence of q-learning: A simple proof. *Institute Of Systems and Robotics, Tech. Rep*, pages 1–4, 2001.

[111] Donald E Kirk. *Optimal control theory: an introduction*. Courier Corporation, 2004.

[112] Richard S Sutton, David A McAllester, Satinder P Singh, Yishay Mansour, et al. Policy gradient methods for reinforcement learning with function approximation. In *NIPs*, volume 99, pages 1057–1063. Citeseer, 1999.

[113] Evan Greensmith, Peter L Bartlett, and Jonathan Baxter. Variance reduction techniques for gradient estimates in reinforcement learning. *Journal of Machine Learning Research*, 5(9), 2004.

[114] Vijay R Konda and John N Tsitsiklis. Actor-critic algorithms. In *Advances in neural information processing systems*, pages 1008–1014, 2000.

[115] Volodymyr Mnih, Adria Puigdomenech Badia, Mehdi Mirza, Alex Graves, Timothy Lillicrap, Tim Harley, David Silver, and Koray Kavukcuoglu. Asynchronous methods for deep reinforcement learning. In *International conference on machine learning*, pages 1928–1937. PMLR, 2016.

[116] Christopher John Cornish Hellaby Watkins. Learning from delayed rewards. 1989.

[117] Long-Ji Lin. *Self-supervised Learning by Reinforcement and Artificial Neural Networks*. PhD thesis, Carnegie Mellon University, School of Computer Science, 1992.

[118] Francisco S Melo and M Isabel Ribeiro. Q-learning with linear function approximation. In *International Conference on Computational Learning Theory*, pages 308–322. Springer, 2007.

[119] Volodymyr Mnih, Koray Kavukcuoglu, David Silver, Andrei A Rusu, Joel Veness, Marc G Bellemare, Alex Graves, Martin Riedmiller, Andreas K Fidjeland, Georg Ostrovski, et al. Human-level control through deep reinforcement learning. *nature*, 518(7540):529–533, 2015.

[120] Marcello Benedetti, Erika Lloyd, Stefan Sack, and Mattia Fiorentini. Parameterized quantum circuits as machine learning models. *Quantum Science and Technology*, 4(4):043001, 2019.

[121] Vedran Dunjko and Peter Wittek. A non-review of quantum machine learning: trends and explorations. *Quantum Views*, 4:32, 2020.

[122] M Cerezo, Guillaume Verdon, Hsin-Yuan Huang, Lukasz Cincio, and Patrick J Coles. Challenges and opportunities in quantum machine learning. *Nature Computational Science*, 2(9):567–576, 2022.

[123] Maria Schuld, Ilya Sinayskiy, and Francesco Petruccione. The quest for a quantum neural network. *Quantum Information Processing*, 13:2567–2586, 2014.

[124] Bernhard Schölkopf, Alexander J Smola, Francis Bach, et al. *Learning with kernels: support vector machines, regularization, optimization, and beyond*. MIT press, 2002.

[125] Corinna Cortes and Vladimir Vapnik. Support-vector networks. *Machine learning*, 20(3):273–297, 1995.

[126] Maria Schuld and Nathan Killoran. Quantum machine learning in feature hilbert spaces. *Physical review letters*, 122(4):040504, 2019.

[127] Maria Schuld. Supervised quantum machine learning models are kernel methods. *arXiv preprint arXiv:2101.11020*, 2021.

[128] Yunchao Liu, Srinivasan Arunachalam, and Kristan Temme. A rigorous and robust quantum speed-up in supervised machine learning. *Nature Physics*, pages 1–5, 2021.

[129] Evan Peters, João Caldeira, Alan Ho, Stefan Leichenauer, Masoud Mohseni, Hartmut Neven, Panagiotis Spentzouris, Doug Strain, and Gabriel N Perdue. Machine learning of high dimensional data on a noisy quantum processor. *npj Quantum Information*, 7(1):1–5, 2021.

[130] Supanut Thanasilp, Samson Wang, Marco Cerezo, and Zoë Holmes. Exponential concentration and untrainability in quantum kernel methods. *arXiv preprint arXiv:2208.11060*, 2022.

[131] Arthur Jacot, Franck Gabriel, and Clément Hongler. Neural tangent kernel: Convergence and generalization in neural networks. *Advances in neural information processing systems*, 31, 2018.

[132] Zeyuan Allen-Zhu, Yuanzhi Li, and Zhao Song. A convergence theory for deep learning via over-parameterization. arxiv e-prints, art. *arXiv preprint arXiv:1811.03962*, 2018.

[133] SS Du, JD Lee, H Li, L Wang, and X Zhai. Gradient descent finds global minima of deep 774 neural networks. *CoRR abs/1811.03804*, 775, 2018.

[134] Sofiene Jerbi, Lukas J Fiderer, Hendrik Poulsen Nautrup, Jonas M Kübler, Hans J Briegel, and Vedran Dunjko. Quantum machine learning beyond kernel methods. *arXiv preprint arXiv:2110.13162*, 2021.

[135] Norihito Shirai, Kenji Kubo, Kosuke Mitarai, and Keisuke Fujii. Quantum tangent kernel. *arXiv preprint arXiv:2111.02951*, 2021.

[136] Junyu Liu, Francesco Tacchino, Jennifer R Glick, Liang Jiang, and Antonio Mezzacapo. Representation learning via quantum neural tangent kernels. *PRX Quantum*, 3(3):030323, 2022.

[137] Edward Farhi and Hartmut Neven. Classification with quantum neural networks on near term processors. *arXiv preprint arXiv:1802.06002*, 2018.

[138] Edward Grant, Marcello Benedetti, Shuxiang Cao, Andrew Hallam, Joshua Lockhart, Vid Stojevic, Andrew G Green, and Simone Severini. Hierarchical quantum classifiers. *npj Quantum Information*, 4(1):1–8, 2018.

[139] Jin-Guo Liu and Lei Wang. Differentiable learning of quantum circuit born machines. *Physical Review A*, 98(6):062324, 2018.

[140] Marcello Benedetti, Delfina Garcia-Pintos, Oscar Perdomo, Vicente Leyton-Ortega, Yunseong Nam, and Alejandro Perdomo-Ortiz. A generative modeling approach for benchmarking and training shallow quantum circuits. *npj Quantum Information*, 5(1):1–9, 2019.

[141] Ian Goodfellow, Jean Pouget-Abadie, Mehdi Mirza, Bing Xu, David Warde-Farley, Sherjil Ozair, Aaron Courville, and Yoshua Bengio. Generative adversarial networks. *Communications of the ACM*, 63(11):139–144, 2020.

[142] Pierre-Luc Dallaire-Demers and Nathan Killoran. Quantum generative adversarial networks. *Physical Review A*, 98(1):012324, 2018.

[143] Christa Zoufal, Aurélien Lucchi, and Stefan Woerner. Quantum generative adversarial networks for learning and loading random distributions. *npj Quantum Information*, 5(1):1–9, 2019.

[144] He-Liang Huang, Yuxuan Du, Ming Gong, Youwei Zhao, Yulin Wu, Chaoyue Wang, Shaowei Li, Futian Liang, Jin Lin, Yu Xu, et al. Experimental quantum generative adversarial networks for image generation. *Physical Review Applied*, 16(2):024051, 2021.

[145] Murphy Yuezhen Niu, Alexander Zlokapa, Michael Broughton, Sergio Boixo, Masoud Mohseni, Vadim Smelyanskyi, and Hartmut Neven. Entangling quantum generative adversarial networks. *Physical Review Letters*, 128(22):220505, 2022.

[146] John J Hopfield. Neural networks and physical systems with emergent collective computational abilities. *Proceedings of the national academy of sciences*, 79(8):2554–2558, 1982.

[147] Nathan Wiebe and Leonard Wossnig. Generative training of quantum boltzmann machines with hidden units. *arXiv preprint arXiv:1905.09902*, 2019.

[148] Christa Zoufal, Aurélien Lucchi, and Stefan Woerner. Variational quantum boltzmann machines. *Quantum Machine Intelligence*, 3(1):1–15, 2021.

[149] Guillaume Verdon, Trevor McCourt, Enxhell Luzhnica, Vikash Singh, Stefan Leichenauer, and Jack Hidary. Quantum graph neural networks. *arXiv preprint arXiv:1909.12264*, 2019.

[150] Sofiene Jerbi, Casper Gyurik, Simon Marshall, Hans Briegel, and Vedran Dunjko. Parametrized quantum policies for reinforcement learning. *Advances in Neural Information Processing Systems*, 34, 2021.

[151] Dániel Nagy, Zsolt Tabi, Péter Hága, Zsófia Kallus, and Zoltán Zimborás. Photonic quantum policy learning in openai gym. In *2021 IEEE International Conference on Quantum Computing and Engineering (QCE)*, pages 123–129. IEEE, 2021.

[152] Nico Meyer, Daniel D Scherer, Axel Plinge, Christopher Mutschler, and Michael J Hartmann. Quantum policy gradient algorithm with optimized action decoding. *arXiv preprint arXiv:2212.06663*, 2022.

[153] Samuel Yen-Chi Chen, Chao-Han Huck Yang, Jun Qi, Pin-Yu Chen, Xiaoli Ma, and Hsi-Sheng Goan. Variational quantum circuits for deep reinforcement learning. *IEEE Access*, 8:141007–141024, 2020.

[154] Owen Lockwood and Mei Si. Reinforcement learning with quantum variational circuit. In *Proceedings of the AAAI Conference on Artificial Intelligence and Interactive Digital Entertainment*, pages 245–251, 2020.

[155] Dirk Heimann, Hans Hohenfeld, Felix Wiebe, and Frank Kirchner. Quantum deep reinforcement learning for robot navigation tasks. *arXiv preprint arXiv:2202.12180*, 2022.

[156] Maja Franz, Lucas Wolf, Maniraman Periyasamy, Christian Ufrecht, Daniel D Scherer, Axel Plinge, Christopher Mutschler, and Wolfgang Mauerer. Uncovering instabilities in variational-quantum deep q-networks. *arXiv preprint arXiv:2202.05195*, 2022.

[157] Qingfeng Lan. Variational quantum soft actor-critic. *arXiv preprint arXiv:2112.11921*, 2021.

[158] Won Joon Yun, Yunseok Kwak, Jae Pyoung Kim, Hyunhee Cho, Soyi Jung, Jihong Park, and Joongheon Kim. Quantum multi-agent reinforcement learning via variational quantum circuit design. *arXiv preprint arXiv:2203.10443*, 2022.

[159] Chen Zhao and Xiao-Shan Gao. Qdnn: Dnn with quantum neural network layers. *arXiv preprint arXiv:1912.12660*, 2019.

[160] Andrea Mari, Thomas R Bromley, Josh Izaac, Maria Schuld, and Nathan Killoran. Transfer learning in hybrid classical-quantum neural networks. *Quantum*, 4:340, 2020.

[161] Maxwell Henderson, Samriddhi Shakya, Shashindra Pradhan, and Tristan Cook. Quanvolutional neural networks: powering image recognition with quantum circuits. *Quantum Machine Intelligence*, 2(1):1–9, 2020.

[162] Tong Dou, Kaiwei Wang, Zhenwei Zhou, Shilu Yan, and Wei Cui. An unsupervised feature learning for quantum-classical convolutional network with applications to fault detection. In *2021 40th Chinese Control Conference (CCC)*, pages 6351–6355. IEEE, 2021.

[163] Samuel Yen-Chi Chen, Chih-Min Huang, Chia-Wei Hsing, and Ying-Jer Kao. An end-to-end trainable hybrid classical-quantum classifier. *Machine Learning: Science and Technology*, 2(4):045021, 2021.

[164] Emanuel Knill, Raymond Laflamme, and Gerald J Milburn. A scheme for efficient quantum computation with linear optics. *nature*, 409(6816):46–52, 2001.

[165] Nathan Wiebe. Key questions for the quantum machine learner to ask themselves. *New Journal of Physics*, 22(9):091001, 2020.

[166] Ryan LaRose and Brian Coyle. Robust data encodings for quantum classifiers. *Physical Review A*, 102(3):032420, 2020.

[167] Francisco Javier Gil Vidal and Dirk Oliver Theis. Input redundancy for parameterized quantum circuits. *Frontiers in Physics*, 8:297, 2020.

[168] Adrián Pérez-Salinas, Alba Cervera-Lierta, Elies Gil-Fuster, and José I Latorre. Data re-uploading for a universal quantum classifier. *Quantum*, 4:226, 2020.

[169] Maria Schuld, Ryan Sweke, and Johannes Jakob Meyer. Effect of data encoding on the expressive power of variational quantum-machine-learning models. *Physical Review A*, 103(3):032430, 2021.

[170] Abhinav Kandala, Antonio Mezzacapo, Kristan Temme, Maika Takita, Markus Brink, Jerry M Chow, and Jay M Gambetta. Hardware-efficient variational quantum eigensolver for small molecules and quantum magnets. *Nature*, 549(7671):242–246, 2017.

[171] Han Zheng, Zimu Li, Junyu Liu, Sergii Strelchuk, and Risi Kondor. Speeding up learning quantum states through group equivariant convolutional quantum ansatze. *arXiv preprint arXiv:2112.07611*, 2021.

[172] Péter Mernyei, Konstantinos Meichanetzidis, and İsmail İlkan Ceylan. Equivariant quantum graph circuits. *arXiv preprint arXiv:2112.05261*, 2021.

[173] Martín Larocca, Frédéric Sauvage, Faris M. Sbahi, Guillaume Verdon, Patrick J. Coles, and M. Cerezo. Group-invariant quantum machine learning. *arXiv preprint arXiv:2205.02261*, 2022.

[174] Andrea Skolik, Michele Cattelan, Sheir Yarkoni, Thomas Bäck, and Vedran Dunjko. Equivariant quantum circuits for learning on weighted graphs. *arXiv preprint arXiv:2205.06109*, 2022.

[175] Johannes Jakob Meyer, Marian Mularski, Elies Gil-Fuster, Antonio Anna Mele, Francesco Arzani, Alissa Wilms, and Jens Eisert. Exploiting symmetry in variational quantum machine learning. *arXiv preprint arXiv:2205.06217*, 2022.

[176] Quynh T Nguyen, Louis Schatzki, Paolo Braccia, Michael Ragone, Patrick J Coles, Frederic Sauvage, Martin Larocca, and M Cerezo. Theory for equivariant quantum neural networks. *arXiv preprint arXiv:2210.08566*, 2022.

[177] Hsin-Yuan Huang, Michael Broughton, Jordan Cotler, Sitan Chen, Jerry Li, Masoud Mohseni, Hartmut Neven, Ryan Babbush, Richard Kueng, John Preskill, et al. Quantum advantage in learning from experiments. *Science*, 376(6598):1182–1186, 2022.

[178] Hsin-Yuan Huang, Michael Broughton, Masoud Mohseni, Ryan Babbush, Sergio Boixo, Hartmut Neven, and Jarrod R McClean. Power of data in quantum machine learning. *Nature communications*, 12(1):1–9, 2021.

[179] Amira Abbas, David Sutter, Christa Zoufal, Aurélien Lucchi, Alessio Figalli, and Stefan Woerner. The power of quantum neural networks. *Nature Computational Science*, 1(6):403–409, 2021.

[180] Edward Farhi and Aram W Harrow. Quantum supremacy through the quantum approximate optimization algorithm. *arXiv preprint arXiv:1602.07674*, 2016.

[181] Ryan Sweke, Jean-Pierre Seifert, Dominik Hangleiter, and Jens Eisert. On the quantum versus classical learnability of discrete distributions. *Quantum*, 5:417, 2021.

[182] Dave Wecker, Matthew B Hastings, and Matthias Troyer. Progress towards practical quantum variational algorithms. *Physical Review A*, 92(4):042303, 2015.

[183] Giacomo Nannicini. Performance of hybrid quantum-classical variational heuristics for combinatorial optimization. *Physical Review E*, 99(1):013304, 2019.

[184] Ken M Nakanishi, Keisuke Fujii, and Synge Todo. Sequential minimal optimization for quantum-classical hybrid algorithms. *Physical Review Research*, 2(4):043158, 2020.

[185] Leo Zhou, Sheng-Tao Wang, Soonwon Choi, Hannes Pichler, and Mikhail D Lukin. Quantum approximate optimization algorithm: performance, mechanism, and implementation on near-term devices. *Preprint at arXiv:1812.01041*, 2018.

[186] Kevin Jeffery Sung, Jiahao Yao, Matthew Harrigan, Nicholas Rubin, Zhang Jiang, Lin Lin, Ryan Babbush, and Jarrod McClean. Using models to improve optimizers for variational quantum algorithms. *Quantum Science and Technology*, 2020.

[187] Emanuel Knill, Gerardo Ortiz, and Rolando D. Somma. Optimal quantum measurements of expectation values of observables. *Physical Review A*, 75(1):012328, jan 2007.

[188] Bobak Toussi Kiani, Seth Lloyd, and Reevu Maity. Learning Unitaries by Gradient Descent. 2020.

[189] Chufan Lyu, Victor Montenegro, and Abolfazl Bayat. Accelerated variational algorithms for digital quantum simulation of many-body ground states. *Quantum*, 4:324, 2020.

[190] Scott E. Fahlman and Christian Lebiere. The cascade-correlation learning architecture. *Advances in Neural Information Processing*, 2:524—-532, 1990.

[191] Chris Hettinger, Tanner Christensen, Ben Ehlert, Jeffrey Humpherys, Tyler Jarvis, and Sean Wade. Forward Thinking: Building and Training Neural Networks One Layer at a Time. In *31st Conference on Neural Information Processing Systems*, 2017.

[192] Geoffrey E. Hinton, Simon Osindero, and Yee-Whye Teh. A Fast Learning Algorithm for Deep Belief Nets. *Neural Computation*, 18(7):1527–1554, jul 2006.

[193] Yoshua Bengio, Yoshua Bengio, Pascal Lamblin, Dan Popovici, and Hugo Larochelle. Greedy layer-wise training of deep networks. *Advances in Neural Information Processing*, 2007.

[194] Mohammad H Amin, Evgeny Andriyash, Jason Rolfe, Bohdan Kulchytskyy, and Roger Melko. Quantum boltzmann machine. *Physical Review X*, 8(2):021050, 2018.

[195] Seth Lloyd and Christian Weedbrook. Quantum generative adversarial learning. *Physical review letters*, 121(4):040502, 2018.

[196] Shouvanik Chakrabarti, Huang Yiming, Tongyang Li, Soheil Feizi, and Xiaodi Wu. Quantum wasserstein generative adversarial networks. In *Advances in Neural Information Processing Systems*, pages 6781–6792, 2019.

[197] A Hamann, V Dunjko, and S Wölk. Quantum-accessible reinforcement learning beyond strictly epochal environments. *arXiv preprint arXiv:2008.01481*, 2020.

[198] Sofiene Jerbi, Lea M Trenkwalder, Hendrik Poulsen Nautrup, Hans J Briegel, and Vedran Dunjko. Quantum enhancements for deep reinforcement learning in large spaces. *PRX Quantum*, 2(1):010328, 2021.

[199] Shaojun Wu, Shan Jin, Dingding Wen, and Xiaoting Wang. Quantum reinforcement learning in continuous action space. *arXiv preprint arXiv:2012.10711*, 2020.

[200] Christopher Berner, Greg Brockman, Brooke Chan, Vicki Cheung, Przemyslaw Debiak, Christy Dennison, David Farhi, Quirin Fischer, Shariq Hashme, Chris Hesse, et al. Dota 2 with large scale deep reinforcement learning. *arXiv preprint arXiv:1912.06680*, 2019.

[201] Oriol Vinyals, Igor Babuschkin, Wojciech M Czarnecki, Michaël Mathieu, Andrew Dudzik, Junyoung Chung, David H Choi, Richard Powell, Timo Ewalds, Petko Georgiev, et al. Grandmaster level in starcraft ii using multi-agent reinforcement learning. *Nature*, 575(7782):350–354, 2019.

[202] Leslie N Smith. A disciplined approach to neural network hyper-parameters: Part 1–learning rate, batch size, momentum, and weight decay. *arXiv preprint arXiv:1803.09820*, 2018.

[203] Ziyu Ye, Andrew Gilman, Qihang Peng, Kelly Levick, Pamela Cosman, and Larry Milstein. Comparison of neural network architectures for spectrum sensing. In *2019 IEEE Globecom Workshops (GC Wkshps)*, pages 1–6. IEEE, 2019.

[204] Hao Yu, Tiantian Xie, Michael Hamilton, and Bogdan Wilamowski. Comparison of different neural network architectures for digit image recognition. In *2011 4th International Conference on Human System Interactions, HSI 2011*, pages 98–103. IEEE, 2011.

[205] F Cordoni. A comparison of modern deep neural network architectures for energy spot price forecasting. *Digital Finance*, 2:189–210, 2020.

[206] Tomasz Szandała. Review and comparison of commonly used activation functions for deep neural networks. In *Bio-inspired Neurocomputing*, pages 203–224. Springer, 2021.

[207] Chigozie Nwankpa, Winifred Ijomah, Anthony Gachagan, and Stephen Marshall. Activation functions: Comparison of trends in practice and research for deep learning. *arXiv preprint arXiv:1811.03378*, 2018.

[208] Sebastian Urban. *Neural network architectures and activation functions: A gaussian process approach.* PhD thesis, Technische Universität München, 2018.

[209] Leslie N Smith. Cyclical learning rates for training neural networks. In *2017 IEEE winter conference on applications of computer vision (WACV)*, pages 464–472. IEEE, 2017.

[210] Thomas Elsken, Jan Hendrik Metzen, and Frank Hutter. Neural architecture search: A survey. *The Journal of Machine Learning Research*, 20(1):1997–2017, 2019.

[211] Frank Hutter, Lars Kotthoff, and Joaquin Vanschoren. *Automated machine learning: methods, systems, challenges*. Springer Nature, 2019.

[212] Bobak Toussi Kiani, Seth Lloyd, and Reevu Maity. Learning unitaries by gradient descent. *arXiv preprint arXiv:2001.11897*, 2020.

[213] Roeland Wiersema, Cunlu Zhou, Yvette de Sereville, Juan Felipe Carrasquilla, Yong Baek Kim, and Henry Yuen. Exploring entanglement and optimization within the hamiltonian variational ansatz. *PRX Quantum*, 1(2):020319, 2020.

[214] Andrea Skolik, Jarrod R McClean, Masoud Mohseni, Patrick van der Smagt, and Martin Leib. Layerwise learning for quantum neural networks. *Quantum Machine Intelligence*, 3 (1):1–11, 2021.

[215] Sukin Sim, Peter D Johnson, and Alán Aspuru-Guzik. Expressibility and entangling capability of parameterized quantum circuits for hybrid quantum-classical algorithms. *Advanced Quantum Technologies*, 2(12):1900070, 2019.

[216] Sukin Sim, Jhonathan Romero Fontalvo, Jérôme F Gonthier, and Alexander A Kunitsa. Adaptive pruning-based optimization of parameterized quantum circuits. *Quantum Science and Technology*, 2021.

[217] Xiaoyuan Liu, Anthony Angone, Ruslan Shaydulin, Ilya Safro, Yuri Alexeev, and Lukasz Cincio. Layer vqe: A variational approach for combinatorial optimization on noisy quantum computers. *arXiv preprint arXiv:2102.05566*, 2021.

[218] OpenAI. Openai gym wiki, cartpole v0 https://github.com/openai/gym/wiki/cartpole-v0, 2021.

[219] Greg Brockman, Vicki Cheung, Ludwig Pettersson, Jonas Schneider, John Schulman, Jie Tang, and Wojciech Zaremba. Openai gym. *arXiv preprint arXiv:1606.01540*, 2016.

[220] Kei Ota, Devesh K Jha, and Asako Kanezaki. Training larger networks for deep reinforcement learning. *arXiv preprint arXiv:2102.07920*, 2021.

[221] Andrea Skolik. Code used in this work https://github.com/askolik/quantum_agents, 2021.

[222] Vedran Dunjko, Yi-Kai Liu, Xingyao Wu, and Jacob M Taylor. Exponential improvements for quantum-accessible reinforcement learning. *arXiv preprint arXiv:1710.11160*, 2017.

[223] OpenAI. Openai gym wiki, frozen lake v0 https://github.com/openai/gym/wiki/frozenlake-v0, 2021.

[224] Michael Broughton, Guillaume Verdon, Trevor McCourt, Antonio J Martinez, Jae Hyeon Yoo, Sergei V Isakov, Philip Massey, Murphy Yuezhen Niu, Ramin Halavati, Evan Peters, et al. Tensorflow quantum: A software framework for quantum machine learning. *arXiv preprint arXiv:2003.02989*, 2020.

[225] Google. Cirq, https://quantumai.google/cirq, 2021.

[226] OpenAI. Openai gym leaderboard https://github.com/openai/gym/wiki/leaderboard, 2021.

[227] Tavis Bennett, Edric Matwiejew, Sam Marsh, and Jingbo B Wang. Quantum walk-based vehicle routing optimisation. *Frontiers in Physics*, page 692, 2021.

[228] Alberto Peruzzo, Jarrod McClean, Peter Shadbolt, Man-Hong Yung, Xiao-Qi Zhou, Peter J Love, Alán Aspuru-Guzik, and Jeremy L O'brien. A variational eigenvalue solver on a photonic quantum processor. *Nature communications*, 5:4213, 2014.

[229] Jarrod R McClean, Sergio Boixo, Vadim N Smelyanskiy, Ryan Babbush, and Hartmut Neven. Barren plateaus in quantum neural network training landscapes. *Nature communications*, 9(1):1–6, 2018.

[230] Patrick van der Smagt and Gerd Hirzinger. Why feed-forward networks are in a bad shape. In *International Conference on Artificial Neural Networks*, pages 159–164. Springer, 1998.

[231] Aditya Ramesh, Mikhail Pavlov, Gabriel Goh, Scott Gray, Chelsea Voss, Alec Radford, Mark Chen, and Ilya Sutskever. Zero-shot text-to-image generation.

In *International Conference on Machine Learning*, pages 8821–8831. PMLR, 2021.

[232] Jie Zhou, Ganqu Cui, Shengding Hu, Zhengyan Zhang, Cheng Yang, Zhiyuan Liu, Lifeng Wang, Changcheng Li, and Maosong Sun. Graph neural networks: A review of methods and applications. *AI Open*, 1:57–81, 2020.

[233] Louis-Paul Henry, Slimane Thabet, Constantin Dalyac, and Loïc Henriet. Quantum evolution kernel: Machine learning on graphs with programmable arrays of qubits. *Physical Review A*, 104(3):032416, 2021.

[234] Jin Zheng, Qing Gao, and Yanxuan Lü. Quantum graph convolutional neural networks. In *2021 40th Chinese Control Conference (CCC)*, pages 6335–6340. IEEE, 2021.

[235] Andrew Lucas. Ising formulations of many np problems. *Frontiers in physics*, 2:5, 2014.

[236] Robin M Schmidt. Recurrent neural networks (rnns): A gentle introduction and overview. *arXiv preprint arXiv:1912.05911*, 2019.

[237] Jianwu Long et al. A graph neural network for superpixel image classification. In *Journal of Physics: Conference Series*, volume 1871, page 012071. IOP Publishing, 2021.

[238] Yanhu Chen, Cen Wang, Hongxiang Guo, et al. Novel architecture of parameterized quantum circuit for graph convolutional network. *arXiv preprint arXiv:2203.03251*, 2022.

[239] Hanjun Dai, Elias B Khalil, Yuyu Zhang, Bistra Dilkina, and Le Song. Learning combinatorial optimization algorithms over graphs. *arXiv preprint arXiv:1704.01665*, 2017.

[240] Wenqi Fan, Yao Ma, Qing Li, Yuan He, Eric Zhao, Jiliang Tang, and Dawei Yin. Graph neural networks for social recommendation. In *The world wide web conference*, pages 417–426, 2019.

[241] Asier Ozaeta, Wim van Dam, and Peter L McMahon. Expectation values from the single-layer quantum approximate optimization algorithm on ising problems. *Quantum Science and Technology*, 2022.

[242] Martin Larocca, Piotr Czarnik, Kunal Sharma, Gopikrishnan Muraleedharan, Patrick J Coles, and M Cerezo. Diagnosing barren plateaus with tools from quantum optimal control. *Quantum*, 6:824, 2022.

[243] Song Mei, Theodor Misiakiewicz, and Andrea Montanari. Learning with invariances in random features and kernel models. In *Conference on Learning Theory*, pages 3351–3418. PMLR, 2021.

[244] Matthias C Caro, Hsin-Yuan Huang, Marco Cerezo, Kunal Sharma, Andrew Sornborger, Lukasz Cincio, and Patrick J Coles. Generalization in quantum machine learning from few training data. *Nature communications*, 13(1):1–11, 2022.

[245] Yoshua Bengio, Andrea Lodi, and Antoine Prouvost. Machine learning for combinatorial optimization: a methodological tour d'horizon. *European Journal of Operational Research*, 290(2):405–421, 2021.

[246] Oriol Vinyals, Meire Fortunato, and Navdeep Jaitly. Pointer networks. *arXiv preprint arXiv:1506.03134*, 2015.

[247] Azalia Mirhoseini, Anna Goldie, Mustafa Yazgan, Joe Jiang, Ebrahim Songhori, Shen Wang, Young-Joon Lee, Eric Johnson, Omkar Pathak, Sungmin Bae, et al. Chip placement with deep reinforcement learning. *arXiv preprint arXiv:2004.10746*, 2020.

[248] Mohammadreza Nazari, Afshin Oroojlooy, Lawrence V Snyder, and Martin Takáč. Reinforcement learning for solving the vehicle routing problem. *arXiv preprint arXiv:1802.04240*, 2018.

[249] Kunal Marwaha. Local classical max-cut algorithm outperforms $p = 2$ qaoa on high-girth regular graphs. *Quantum*, 5:437, 2021.

[250] Mario Szegedy. What do qaoa energies reveal about graphs? *arXiv preprint arXiv:1912.12277*, 2019.

[251] Gilbert H Harman, Sanjeev R Kulkarni, and Hariharan Narayanan. $\sin(\omega x)$ can approximate almost every finite set of samples. *Constructive Approximation*, 42(2):303–311, 2015.

[252] Mauro ES Morales, Jacob D Biamonte, and Zoltán Zimborás. On the universality of the quantum approximate optimization algorithm. *Quantum Information Processing*, 19(9):1–26, 2020.

[253] Python-TSP.

[254] Nicos Christofides. Worst-case analysis of a new heuristic for the travelling salesman problem. Technical report, Carnegie-Mellon Univ Pittsburgh Pa Management Sciences Research Group, 1976.

[255] Leo Zhou, Sheng-Tao Wang, Soonwon Choi, Hannes Pichler, and Mikhail D Lukin. Quantum approximate optimization algorithm: Performance, mechanism, and implementation on near-term devices. *Physical Review X*, 11(2):021067, 2020.

[256] Adam Glos, Alexandra Krawiec, and Zoltán Zimborás. Space-efficient binary optimization for variational computing. *arXiv preprint arXiv:2009.07309*, 2020.

[257] Edward Farhi, David Gamarnik, and Sam Gutmann. The quantum approximate optimization algorithm needs to see the whole graph: A typical case. *arXiv preprint arXiv:2004.09002*, 2020.

[258] Sergey Bravyi, Alexander Kliesch, Robert Koenig, and Eugene Tang. Obstacles to variational quantum optimization from symmetry protection. *Physical review letters*, 125(26):260505, 2020.

[259] Leo Zhou, Sheng-Tao Wang, Soonwon Choi, Hannes Pichler, and Mikhail D Lukin. Quantum approximate optimization algorithm: Performance, mechanism, and implementation on near-term devices. *Physical Review X*, 10(2):021067, 2020.

[260] Ruslan Shaydulin, Phillip C Lotshaw, Jeffrey Larson, James Ostrowski, and Travis S Humble. Parameter transfer for quantum approximate optimization of weighted maxcut. *arXiv preprint arXiv:2201.11785*, 2022.

[261] Laura Gentini, Alessandro Cuccoli, Stefano Pirandola, Paola Verrucchi, and Leonardo Banchi. Noise-resilient variational hybrid quantum-classical optimization. *Physical Review A*, 102(5):052414, 2020.

[262] Kam-Chuen Jim, C Lee Giles, and Bill G Horne. An analysis of noise in recurrent neural networks: convergence and generalization. *IEEE Transactions on neural networks*, 7(6):1424–1438, 1996.

[263] Hyeonwoo Noh, Tackgeun You, Jonghwan Mun, and Bohyung Han. Regularizing deep neural networks by noise: Its interpretation and optimization. *Advances in Neural Information Processing Systems*, 30, 2017.

[264] Alex Graves. Practical variational inference for neural networks. *Advances in neural information processing systems*, 24, 2011.

[265] Alex Graves, Abdel-rahman Mohamed, and Geoffrey Hinton. Speech recognition with deep recurrent neural networks. In *2013 IEEE international conference on acoustics, speech and signal processing*, pages 6645–6649. Ieee, 2013.

[266] Emilio Rafael Balda, Arash Behboodi, and Rudolf Mathar. Adversarial examples in deep neural networks: An overview. *Deep Learning: Algorithms and Applications*, pages 31–65, 2020.

[267] Cihang Xie, Jianyu Wang, Zhishuai Zhang, Zhou Ren, and Alan Yuille. Mitigating adversarial effects through randomization. *arXiv preprint arXiv:1711.01991*, 2017.

[268] Justin Gilmer, Nicolas Ford, Nicholas Carlini, and Ekin Cubuk. Adversarial examples are a natural consequence of test error in noise. In *International Conference on Machine Learning*, pages 2280–2289. PMLR, 2019.

[269] Florian Jaeckle and M Pawan Kumar. Generating adversarial examples with graph neural networks. In *Uncertainty in Artificial Intelligence*, pages 1556–1564. PMLR, 2021.

[270] Ian J Goodfellow, Jonathon Shlens, and Christian Szegedy. Explaining and harnessing adversarial examples. *arXiv preprint arXiv:1412.6572*, 2014.

[271] Nitish Srivastava, Geoffrey Hinton, Alex Krizhevsky, Ilya Sutskever, and Ruslan Salakhutdinov. Dropout: a simple way to prevent neural networks from overfitting. *The journal of machine learning research*, 15(1):1929–1958, 2014.

[272] Jinfeng Zeng, Zipeng Wu, Chenfeng Cao, Chao Zhang, Shi-Yao Hou, Pengxiang Xu, and Bei Zeng. Simulating noisy variational quantum eigensolver with local noise models. *Quantum Engineering*, 3(4):e77, 2021.

[273] Mahabubul Alam, Abdullah Ash-Saki, and Swaroop Ghosh. Analysis of quantum approximate optimization algorithm under realistic noise in superconducting qubits. *arXiv preprint arXiv:1907.09631*, 2019.

[274] Junyu Liu, Frederik Wilde, Antonio Anna Mele, Liang Jiang, and Jens Eisert. Noise can be helpful for variational quantum algorithms. *arXiv preprint arXiv:2210.06723*, 2022.

[275] Jingkang Wang, Yang Liu, and Bo Li. Reinforcement learning with perturbed rewards. In *Proceedings of the AAAI conference on artificial intelligence*, volume 34, pages 6202–6209, 2020.

[276] Sandy Huang, Nicolas Papernot, Ian Goodfellow, Yan Duan, and Pieter Abbeel. Adversarial attacks on neural network policies. *arXiv preprint arXiv:1702.02284*, 2017.

[277] Jernej Kos and Dawn Song. Delving into adversarial attacks on deep policies. *arXiv preprint arXiv:1705.06452*, 2017.

[278] Yang Yu. Towards sample efficient reinforcement learning. In *IJCAI*, pages 5739–5743, 2018.

[279] Owen Lockwood and Mei Si. Playing atari with hybrid quantum-classical reinforcement learning. In *NeurIPS 2020 Workshop on Pre-registration in Machine Learning*, pages 285–301. PMLR, 2021.

[280] Kosuke Ito, Wataru Mizukami, and Keisuke Fujii. Universal noise-precision relations in variational quantum algorithms. *arXiv preprint arXiv:2106.03390*, 2021.

[281] Andrea Skolik and Stefano Mangini. Code that was used for training of noisy quantum agents. https://github.com/askolik/noisy_qrl, 2022.

[282] OpenAI. Openai gym. https://github.com/openai/gym/wiki, 2022. Accessed: 06-09-2022.

[283] Tensorflow quantum rl tutorial. https://www.tensorflow.org/quantum/tutorials/quantum_reinforcement_learning. Accessed: 06-09-2022.

[284] Aleksandrs Slivkins et al. Introduction to multi-armed bandits. *Foundations and Trends® in Machine Learning*, 12(1-2):1–286, 2019.

[285] Tze Leung Lai, Herbert Robbins, et al. Asymptotically efficient adaptive allocation rules. *Advances in applied mathematics*, 6(1):4–22, 1985.

[286] Peter Auer. Using confidence bounds for exploitation-exploration trade-offs. *Journal of Machine Learning Research*, 3(Nov):397–422, 2002.

[287] Zhenyu Cai, Xiaosi Xu, and Simon C Benjamin. Mitigating coherent noise using pauli conjugation. *npj Quantum Information*, 6(1):1–9, 2020.

[288] Hsin-Yuan Huang, Richard Kueng, and John Preskill. Predicting many properties of a quantum system from very few measurements. *Nature Physics*, 16(10):1050–1057, 2020.

[289] Zbigniew Puchała and Jarosław Adam Miszczak. Symbolic integration with respect to the haar measure on the unitary group. *arXiv preprint arXiv:1109.4244*, 2011.

[290] Hado Van Hasselt, Arthur Guez, and David Silver. Deep reinforcement learning with double q-learning. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 30, 2016.

[291] Google. Documentation of depolarizing channel in cirq. https://quantumai.google/reference/python/cirq/depolarize, 2022.

[292] Sergei V. Isakov, Dvir Kafri, Orion Martin, Catherine Vollgraff Heidweiller, Wojciech Mruczkiewicz, Matthew P. Harrigan, Nicholas C. Rubin, Ross Thomson, Michael Broughton, Kevin Kissell, Evan Peters, Erik Gustafson, Andy C. Y. Li, Henry Lamm, Gabriel Perdue, Alan K. Ho, Doug Strain, and Sergio Boixo. Simulations of quantum circuits with approximate noise using qsim and cirq, 2021.

[293] Michael A. Nielsen and Isaac L. Chuang. *Quantum computation and quantum information*. Cambridge University Press, Cambridge, UK, 2010.

[294] Timothy Proctor, Stefan Seritan, Kenneth Rudinger, Erik Nielsen, Robin Blume-Kohout, and Kevin Young. Scalable randomized benchmarking of quantum computers using mirror circuits. *Physical Review Letters*, 129(15):150502, 2022.

[295] Joseph Vovrosh, Kiran E Khosla, Sean Greenaway, Christopher Self, Myungshik S Kim, and Johannes Knolle. Simple mitigation of global depolarizing errors in quantum simulations. *Physical Review E*, 104(3):035309, 2021.

[296] Easwar Magesan, Jay M Gambetta, and Joseph Emerson. Characterizing quantum gates via randomized benchmarking. *Physical Review A*, 85(4):042311, 2012.

[297] C Ryan-Anderson, NC Brown, MS Allman, B Arkin, G Asa-Attuah, C Baldwin, J Berg, JG Bohnet, S Braxton, N Burdick, et al. Implementing fault-tolerant entangling gates on the five-qubit code and the color code. *arXiv preprint arXiv:2208.01863*, 2022.

[298] IBM. Ibmquantum. https://quantum-computing.ibm.com/, 2022.

[299] Elijah Pelofske, Andreas Bärtschi, and Stephan Eidenbenz. Quantum volume in practice: What users can expect from nisq devices. *arXiv preprint arXiv:2203.03816*, 2022.

[300] IBM Quantum Experience. IBM Quantum Experience. https://quantum-computing.ibm.com/services/resources?tab=systems, 2022.

[301] Xun Gao and Luming Duan. Efficient classical simulation of noisy quantum computation. *arXiv preprint arXiv:1810.03176*, 2018.

[302] Ryan LaRose, Andrea Mari, Sarah Kaiser, Peter J Karalekas, Andre A Alves, Piotr Czarnik, Mohamed El Mandouh, Max H Gordon, Yousef Hindy, Aaron Robertson, et al. Mitiq: A software package for error mitigation on noisy quantum computers. *Quantum*, 6:774, 2022.

[303] Vincent Russo, Andrea Mari, Nathan Shammah, Ryan LaRose, and William J Zeng. Testing platform-independent quantum error mitigation on noisy quantum computers. *arXiv preprint arXiv:2210.07194*, 2022.

[304] Samson Wang, Piotr Czarnik, Andrew Arrasmith, Marco Cerezo, Lukasz Cincio, and Patrick J Coles. Can error mitigation improve trainability of noisy variational quantum algorithms? *arXiv preprint arXiv:2109.01051*, 2021.

[305] Gian-Carlo Wick. The evaluation of the collision matrix. *Physical review*, 80(2):268, 1950.

[306] Patrick Huembeli and Alexandre Dauphin. Characterizing the loss landscape of variational quantum circuits. *Quantum Science and Technology*, 6(2):025011, 2021.

[307] Motohisa Fukuda, Robert König, and Ion Nechita. Rtni—a symbolic integrator for haar-random tensor networks. *Journal of Physics A: Mathematical and Theoretical*, 52(42):425303, 2019.

[308] Robert W. Keener. *Theoretical Statistics: Topics for a Core Course*. Springer Texts in Statistics. Springer, 1 edition, 2010.