



Universiteit  
Leiden  
The Netherlands

## Resource frugal optimizer for quantum machine learning

MOUSSA, C.; Gordon, M.H.; Baczyk, M.; Cerezo, M.; Cincio, L.; Coles, P.J.

### Citation

MOUSSA, C., Gordon, M. H., Baczyk, M., Cerezo, M., Cincio, L., & Coles, P. J. (2023). Resource frugal optimizer for quantum machine learning. *Arxiv*. Retrieved from <https://hdl.handle.net/1887/3566554>

Version: Not Applicable (or Unknown)

License: [Leiden University Non-exclusive license](#)

Downloaded from: <https://hdl.handle.net/1887/3566554>

**Note:** To cite this publication please use the final published version (if applicable).

# Resource frugal optimizer for quantum machine learning

Charles Moussa,<sup>1,2,\*</sup> Max Hunter Gordon,<sup>1,3</sup> Michal Baczyk,<sup>1,4</sup> M. Cerezo,<sup>5</sup> Lukasz Cincio,<sup>1</sup> and Patrick J. Coles<sup>1</sup>

<sup>1</sup>*Theoretical Division, Los Alamos National Laboratory, Los Alamos, New Mexico 87545, USA*

<sup>2</sup>*LIACS, Leiden University, Niels Bohrweg 1, 2333 CA Leiden, Netherlands*

<sup>3</sup>*Instituto de Física Teórica, UAM/CSIC, Universidad Autónoma de Madrid, Madrid 28049, Spain*

<sup>4</sup>*Faculty of Physics, University of Warsaw, ulica Pasteura 5, 02-093 Warsaw, Poland*

<sup>5</sup>*Information Sciences, Los Alamos National Laboratory, Los Alamos, NM 87545, USA*

Quantum-enhanced data science, also known as quantum machine learning (QML), is of growing interest as an application of near-term quantum computers. Variational QML algorithms have the potential to solve practical problems on real hardware, particularly when involving quantum data. However, training these algorithms can be challenging and calls for tailored optimization procedures. Specifically, QML applications can require a large shot-count overhead due to the large datasets involved. In this work, we advocate for simultaneous random sampling over both the dataset as well as the measurement operators that define the loss function. We consider a highly general loss function that encompasses many QML applications, and we show how to construct an unbiased estimator of its gradient. This allows us to propose a shot-frugal gradient descent optimizer called Refoqus (REsource Frugal Optimizer for QUantum Stochastic gradient descent). Our numerics indicate that Refoqus can save several orders of magnitude in shot cost, even relative to optimizers that sample over measurement operators alone.

## I. INTRODUCTION

A new kind of data is emerging in recent times: quantum data. Tabletop quantum experiments and analog quantum simulators produce interesting sets of quantum states that must be characterized. Moreover, the rise of digital quantum computers is leading to the discovery of novel quantum circuits that can, once again, produce quantum states of interest. Quantum sensing, quantum phase diagrams, quantum error correction, and quantum dynamics are some of the areas that stand to benefit from quantum data analysis.

Classical machine learning was developed for the processing of classical data, but it is necessarily inefficient at processing quantum data. This issue has given rise to the field of quantum machine learning (QML) [1, 2]. QML has seen the proposal of parameterized quantum models, such as quantum neural networks [3–6], that could efficiently process quantum data. Variational QML, which involves classically training a parameterized quantum model, is indeed a leading candidate for implementing QML in the near term.

Variational QML, which we will henceforth refer to as QML for simplicity, has faced various sorts of trainability issues. Exponentially vanishing gradients, known as barren plateaus [7–15], as well as the prevalence of local minima [16, 17] are two issues that can impact the complexity of the training process. Quantum hardware noise also impacts trainability [18, 19]. All of these issues contribute to increasing the number of shots and iterations required to minimize the QML loss function. Indeed, a detailed shot-cost analysis has painted a concerning picture [20].

It is therefore clear that QML requires careful frugality in terms of the resources expended during the optimization process. Indeed, novel optimizers have been developed in response to these challenges. Quantum-aware optimizers aim to replace off-the-shelf classical optimizers with ones that are specifically tailored to the quantum setting [21–23]. Shot-frugal optimizers [24–26] have been proposed in the context of variational quantum eigensolver (VQE), whereby one can sample over terms in the Hamiltonian instead of measuring every term [27]. While significant progress has been made on such optimizers, particularly for VQE, we argue that very little work has been done to specifically tailor optimizers to the QML setting. The cost functions in QML go well beyond those used in VQE and hence QML requires more general tools.

In this work, we generalize previous shot-frugal and iteration-frugal optimizers, such as those in Refs. [24, 25, 27], by extending them to the QML setting. Specifically, we allow for random, weighted sampling over both the input and the output of the loss function estimation circuit. In other words, and as shown in Fig. 1, we allow for sampling over the dataset as well as over the measurement operators used to define the loss function. Our sampling approach allows us to unlock the frugality (i.e., to achieve the full potential) of adaptive stochastic gradient descent optimizers, such as the recently developed iCANS [24] and gCANS [25].

We discuss how our approach applies to various QML applications such as perceptron-based quantum neural networks [11, 28], quantum autoencoders [29], variational quantum principal component analysis (PCA) [30, 31], and classifiers that employ the mean-squared-error loss function [26, 32]. Each of these applications can be unified under one umbrella by considering a generic loss function with a highly general form. Thus we state our main results for this generic loss function. We establish an unbiased estimator for this loss function and its

---

\* [c.moussa@liacs.leidenuniv.nl](mailto:c.moussa@liacs.leidenuniv.nl)

gradient. In turn, this allows us to provide convergence guarantees for certain optimization routines, like stochastic gradient descent. Furthermore, we show that for this general loss function one can use the form of the estimator to inform a strategy that distributes shots to obtain the best shot frugal estimates. We also show how to construct an unbiased estimator of the log likelihood loss function, which can be used in gradient-free shot frugal optimization.

Finally, we numerically investigate the performance of our new optimization approach, which we call Refoqus (REsource Frugal Optimizer for QUantum Stochastic gradient descent). For a quantum PCA task, Refoqus significantly outperforms state-of-the-art optimizers in terms of the shot resources required. Refoqus even outperforms Rosalin [27] - a shot-frugal optimizer that samples only over measurement operators. Hence, Refoqus will be a crucial tool to minimize the number of shots and iterations required in near-term QML implementations.

## II. BACKGROUND

### A. Stochastic Gradient Descent

One of the most popular optimization approaches is gradient descent, which involves the following update rule for the parameter vector:

$$\boldsymbol{\theta}^{(t+1)} = \boldsymbol{\theta}^{(t)} - \alpha \nabla \mathcal{L}(\boldsymbol{\theta}^{(t)}). \quad (1)$$

Here,  $\mathcal{L}$  is the loss function,  $\alpha$  is the learning rate, and  $\boldsymbol{\theta}^{(t)}$  is the parameter vector at iteration  $t$ .

Oftentimes one only has access to noisy estimates of the gradient  $\nabla \mathcal{L}$ , in which case the optimizer is called stochastic gradient descent. In the quantum setting, this situation arises due to shot noise or noise due to sampling from terms in some expansion of the gradient.

### B. Parameter Shift Rule

Estimating the gradient is clearly an essential step in stochastic gradient descent. For this purpose, one useful tool that is often employed in the quantum case is the so-called parameter shift rule [33, 34]. We emphasize that several assumptions go into this rule. Specifically, suppose we assume that the quantum circuit ansatz  $U(\boldsymbol{\theta})$  can be expressed as  $U(\boldsymbol{\theta}) = \prod_x e^{-i\theta_x \sigma_x} W_x$  where  $W_x$  are unparametrized unitaries, and where  $\sigma_x$  are Pauli operators. Moreover, we consider the case when the loss function has the simple form  $\mathcal{L}(\boldsymbol{\theta}) = \langle 0|U^\dagger(\boldsymbol{\theta})HU(\boldsymbol{\theta})|0\rangle$  for some Hermitian operator  $H$ . (Note that we will consider more complicated loss functions in this work, and hence we are just stating this special case for background information.) In this case, the parameter shift rule gives

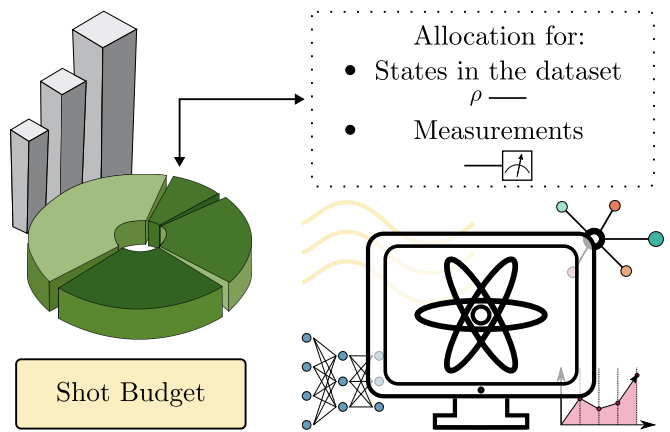


FIG. 1. **Schematic illustration of Refoqus.** Measurements, or shots, are a precious (and expensive) resource in quantum computing. As such, they should be used sparingly and only when absolutely necessary. This is particularly important in variational QML methods where training a model requires continuously calling a quantum device to estimate the loss function or its gradients. Refoqus provides a shot-frugal optimization paradigm where shots are allocated by sampling over the input data and the measurement operators.

$$\partial_x \mathcal{L}(\boldsymbol{\theta}) := \frac{\partial \mathcal{L}(\boldsymbol{\theta})}{\partial \theta_x} = \frac{\mathcal{L}(\boldsymbol{\theta} + \frac{\pi}{2} \boldsymbol{\delta}_x) - \mathcal{L}(\boldsymbol{\theta} - \frac{\pi}{2} \boldsymbol{\delta}_x)}{2}, \quad (2)$$

where  $\boldsymbol{\delta}_x$  is a unit vector with a one on the  $x$ -th component. Equation (2) allows one to estimate the gradient by estimating the loss function as specific points on the landscape. Hence it simplifies the procedure to estimate the gradient.

### C. Lipschitz continuity

The loss function  $\mathcal{L}$  is called Lipschitz continuous if there is some Lipschitz constant  $L \geq 0$  that satisfies the bound

$$\|\nabla \mathcal{L}(\boldsymbol{\theta}_a) - \nabla \mathcal{L}(\boldsymbol{\theta}_b)\| \leq L \|\boldsymbol{\theta}_a - \boldsymbol{\theta}_b\|, \quad (3)$$

for all  $\boldsymbol{\theta}_a, \boldsymbol{\theta}_b \in \text{dom}(\mathcal{L})$ , where  $\|\cdot\|$  is the  $\ell_2$  norm. This property provides a recipe for choosing an appropriate learning rate,  $\alpha$ . Specifically, if (3) holds and we have access to the exact gradient, then choosing  $\alpha \leq 2/L$  is sufficient to guarantee convergence using the update rule in (1).

### D. iCANS

Inspired by an adaptive batch size optimizer used in classical machine learning [35], the iCANS (individual coupled adaptive number of shots) optimizer [24] was introduced as an adaptive method for stochastic gradient in

the context of the variational quantum eigensolver. It allows the number of shots per partial derivative (i.e., gradient component) to vary individually, hence the name iCANS.

Consider the gain (i.e., the decrease in the loss function), denoted  $\mathcal{G}_x$ , associated with updating the  $x$ -th parameter  $\theta_x$ . The goal of iCANS is to maximize the expected gain per shot. That is, for each individual partial derivative, we maximize the shot efficiency:

$$\gamma_x := \frac{\mathbb{E}[\mathcal{G}_x]}{s_x}; \quad x = 1, \dots, d, \quad (4)$$

where  $s_x$  is the shot allocation for gradient component  $x$  and  $d$  is the number of gradient components. Solving for the optimal shot allocation gives:

$$s_x = \frac{2L\alpha}{2 - L\alpha} \frac{\sigma_x^2}{g_x^2}. \quad (5)$$

Here,  $g_x$  is an unbiased estimator for the  $x$ -th gradient component, and  $\sigma_x$  is the standard deviation of a random variable  $X_x$  whose sample mean is  $g_x$ . While iCANS often heuristically outperforms other methods, it can have instabilities.

### E. gCANS

Recently, a potential improvement over iCANS was introduced called gCANS (global coupled adaptive number of shots) [25]. gCANS considers the expected gain  $\mathbb{E}[\mathcal{G}]$  over the entire gradient vector. Then the goal is to maximize the shot efficiency

$$\gamma := \frac{\mathbb{E}[\mathcal{G}]}{\sum_{x=1}^d s_x}, \quad (6)$$

where the sum  $\sum_{x=1}^d s_x$  goes over all components of the gradient. Solving for the optimal shot count then gives:

$$s_x = \frac{2L\alpha}{2 - L\alpha} \frac{\sigma_x \sum_{x'=1}^d \sigma_{x'}}{\|\nabla \mathcal{L}(\boldsymbol{\theta})\|^2}. \quad (7)$$

(Note that an exponential moving average is used to estimate  $\sigma_x$  and  $\|\nabla \mathcal{L}(\boldsymbol{\theta})\|^2$  as their true value is not accessible.) It was proven that gCANS achieves geometric convergence to the optimum, often reducing the number of shots spent for comparable solutions against its predecessor iCANS.

### F. Rosalin

Shot-frugal optimizers like iCANS and gCANS rely on having an unbiased estimator for the gradient or its components. However, this typically places a hard floor on how many shots must be allocated at each iteration, i.e.,

going below this floor could result in a biased estimator. This is because the measurement operator  $H$  is typically composed of multiple non-commuting terms, each of which must be measured individually. Each of these terms must receive some shot allocation to avoid having a biased estimator. However, having this hard floor on the shot requirement is antithetical to the shot-frugal nature of iCANS and gCANS, and ultimately it handicaps these optimizers' ability to achieve shot frugality.

This issue inspired a recent proposal called Rosalin (Random Operator Sampling for Adaptive Learning with Individual Number of shots) [27]. Rosalin employs weighted random sampling of operators in the measurement operator  $H = \sum_j c_j H_j$ , which allows one to achieve an unbiased estimator without a hard floor on the shot requirement. (Even a single shot, provided that it is randomly allocated according to an appropriate probability distribution, can lead to an unbiased estimator.) When combined with the shot allocation methods from iCANS or gCANS, the operator sampling methods in Rosalin were shown to be extremely powerful in the context of molecular chemistry problems, which often have a large number of terms in  $H$ .

We remark that Ref. [27] considered several sampling strategies. Given a budget of  $s_{\text{tot}}$  shots and  $N$  terms, a simple strategy is to distribute shots per term equally ( $s_j = s_{\text{tot}}/N$ ) - referred as uniform deterministic sampling (UDS). Defining  $M = \sum_j |c_j|$ , one can also use weighted deterministic sampling (WDS) where the shots are proportionally distributed:  $s_j = s_{\text{tot}} * \frac{|c_j|}{M}$ . One can add randomness by using  $p_j = \frac{|c_j|}{M}$  to define a (non-uniform) probability distribution to select which term should be measured. This is referred to as weighted random sampling (WRS). Finally, there exists a hybrid approach where one combines WDS with WRS - referred to as weighted hybrid sampling (WHS). Ref. [27] found that the WRS and WHS strategies performed similarly and they both significantly outperformed the UDS and WDS strategies on molecular ground state problems. Because of these results, we choose to focus on the WRS strategy in our work here.

While Rosalin was designed for chemistry problems, it was not designed for QML, where the number of terms in  $H$  is not the only consideration. As discussed below, QML problems involve a (potentially large) dataset of input states. Each input state requires a separate quantum circuit, and hence we are back to the situation of having a hard floor on the shots required due to these multiple input states. This ultimately provides the motivation for our work, which can be viewed as a generalization of Rosalin to the setting of QML.




FIG. 2. **Illustration of a generic variational QML framework.** Given a dataset of quantum states  $\mathcal{S}$ , the input states to the QML model are tensor product states of the form  $\rho_i = \rho_{i_1} \otimes \dots \otimes \rho_{i_m}$ . The number of samples  $m$  depends on the QML task at hand. The parameterized quantum model denoted  $\mathcal{M}_\theta$ , acts on  $m$  copies of the input Hilbert space. Finally, an operator  $H_i$  is measured to estimate a quantity to be used when evaluating a loss function  $\mathcal{L}(\theta)$ . The latter evaluation is then inputted to a classical optimizer that proposes new parameters  $\theta$  in order to minimize the loss. Hence, one can repeat the quantum-classical loop of evaluations and updates until the desired stopping criteria are satisfied.

### III. FRAMEWORK

#### A. Generic Variational QML Framework

Let us present our general framework for discussing (variational) QML methods; see Fig. 2 for an illustration. This framework is meant to unify multiple literature QML algorithms under one umbrella. We discuss how specific literature algorithms are special cases of this framework in Section III B.

In a generic QML setting, one has a training dataset composed of quantum states:

$$\mathcal{S} = \{\rho_i\}_{i=1}^N, \quad (8)$$

where each  $\rho_i$  is a trace-one positive semi-definite matrix, i.e., a density matrix. Each of these training states may come with an associated probability, with the associated probability distribution denoted as

$$\mathcal{P} = \{p_i\}_{i=1}^N. \quad (9)$$

In variational QML, one trains a parameterized quantum model, which we write as  $\mathcal{M}_\theta$  for some set of parameters  $\theta$ . With a large degree of generality, we can assume that  $\mathcal{M}_\theta$  is a linear, completely-positive map. In general,  $\mathcal{M}_\theta$  could act on multiple copies ( $m$  copies) of the input Hilbert space. (Multiple copies allow for non-linear operations on the dataset, which can be important in certain classification tasks.) Hence we allow for the output of this action to be of the form:

$$\mathcal{M}_\theta(\rho_i) = \mathcal{M}_\theta(\rho_{i_1} \otimes \dots \otimes \rho_{i_m}) = \mathcal{M}_\theta \left( \bigotimes_{\alpha=1}^m \rho_{i_\alpha} \right), \quad (10)$$

where we employ the notation  $\rho_i := \rho_{i_1} \otimes \dots \otimes \rho_{i_m}$ . Given that we are allowing for multiple copies of the input space, one can define an effective dataset  $\mathcal{S}_m$  composed of the tensor product of  $m$  states in  $\mathcal{S}$  and an effective probability distribution  $\mathcal{P}_m$ .

A QML loss function is then defined in an operational manner so that it could be estimated on a quantum device. This involves considering the aforementioned mathematical objects as well as a measurement operator, or a

set of measurement operators. We allow for the measurement operator to be tailored to the input state. Hence we write  $H_i$ , with  $\mathbf{i} = \{i_1, \dots, i_m\}$ , as the measurement operator when the input state on  $m$  copies of the Hilbert space is  $\rho_i = \rho_{i_1} \otimes \dots \otimes \rho_{i_m}$ . Moreover, each measurement operator can be decomposed into a linear combination of Hermitian matrices that can be directly measured:

$$H_i = \sum_{j=1}^{t_i} c_{i,j} h_{i,j}. \quad (11)$$

Generically, we could write the loss function as an average over training states chosen from the effective dataset  $D_m$ :

$$\mathcal{L}(\theta) = \sum_i p_i \ell(E_i(\theta)). \quad (12)$$

Here,  $\ell$  is an application-dependent function whose input is a measurable expectation value  $E_i(\theta)$ . Specifically, this expectation value is associated with the  $\rho_i$  input state, with the form

$$E_i(\theta) = \text{Tr}[\mathcal{M}_\theta(\rho_i) H_i]. \quad (13)$$

There are many possible forms for the function  $\ell$ . However, there are multiple QML proposals in the literature that involve a simple linear form:

$$\ell(E_i(\theta)) = E_i(\theta), \quad (14)$$

and in this case, we refer to the overall loss function as a “linear loss function”. Alternatively, non-linear functions are also possible, and we also consider polynomial functions of the form:

$$\ell(E_i(\theta)) = \sum_{z=0}^D a_z [E_i(\theta)]^z, \quad (15)$$

where  $D$  is the degree of the polynomial. In this case, we refer to the loss as a “polynomial loss function”.

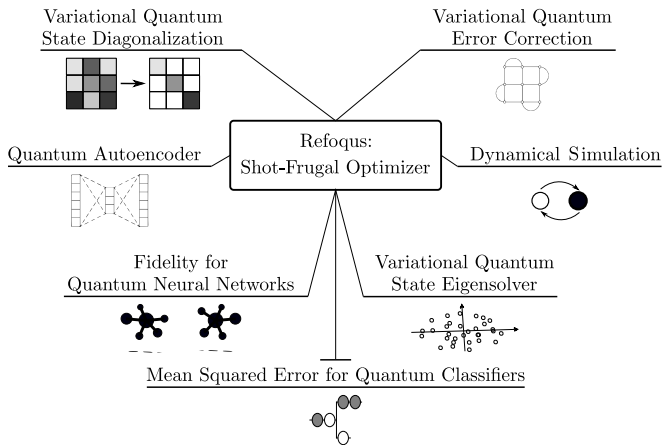


FIG. 3. **Applications benefiting from Refoqus.** Many variational quantum algorithms employ training data, and many QML models (which naturally analyze datasets) are variational in nature. We give examples of both of these cases in this figure. Our Refoqus optimizer is relevant in all cases.

## B. Examples of QML loss functions

Now let us illustrate how various QML loss functions proposed in the literature fall under the previous framework. Crucially, as shown in Fig. (3) these loss functions can be used for a wide range of QML tasks.

### 1. Variational quantum error correction

Variational quantum algorithms to learn device-tailored quantum error correction codes were discussed in Refs. [4, 36]. The loss function involves evaluating the input-output fidelity of the code, averaged over a set of input states. The input state is fed into the encoder  $\mathcal{E}_{\theta_1}$ , then the noise channel  $\mathcal{N}$  acts, followed by the decoder  $\mathcal{D}_{\theta_2}$ . The concatenation of these three channels can be viewed as the overall channel

$$\mathcal{M}_{\theta} = \mathcal{D}_{\theta_2} \circ \mathcal{N} \circ \mathcal{E}_{\theta_1}, \quad (16)$$

with parameter vector  $\theta = (\theta_1, \theta_2)$ . Then the loss function is given by:

$$\mathcal{L}(\theta) = \sum_{|\psi_i\rangle \in \mathcal{S}} \frac{1}{|\mathcal{S}|} \text{Tr}[|\psi_i\rangle\langle\psi_i| \mathcal{M}_{\theta}(|\psi_i\rangle\langle\psi_i|)], \quad (17)$$

where  $\mathcal{S}$  is some appropriately chosen set of states. It is clear that this loss function is of the form in (12) with  $\ell$  having the linear form in (14).

### 2. Quantum autoencoder

Inspired by the success of classical autoencoders, quantum autoencoders [29] were proposed to compress quantum data by reducing the number of qubits needed to

represent the dataset. Consider a bipartite quantum system  $AB$  composed of  $n_A$  and  $n_B$  qubits, respectively, and let  $\{p_i, |\psi_i\rangle\}$  be an ensemble of pure states on  $AB$ . The quantum autoencoder trains a gate sequence  $U(\theta)$  to compress this ensemble into the  $A$  subsystem, such that one can recover each state  $|\psi_i\rangle$  with high fidelity from the information in subsystem  $A$ . One can think of  $B$  as the “trash” since it is discarded after the action of  $U(\theta)$ .

The original proposal [29] employed a loss function that quantified the overlap of the trash state with a fixed pure state:

$$\mathcal{L}_G(\theta) = 1 - \text{Tr}_B[|\mathbf{0}\rangle\langle\mathbf{0}| \rho_B^{\text{out}}] \quad (18)$$

$$= \text{Tr}_{AB}[H_G U(\theta) \rho_{AB}^{\text{in}} U(\theta)^\dagger] \quad (19)$$

$$= \sum_i p_i \text{Tr}_{AB}[H_G U(\theta) |\psi_i\rangle\langle\psi_i| U(\theta)^\dagger]. \quad (20)$$

Here,  $\rho_{AB}^{\text{in}} = \sum_i p_i |\psi_i\rangle\langle\psi_i|$  is the ensemble-average input state,  $\rho_B^{\text{out}} = \text{Tr}_A[U(\theta) \rho_{AB}^{\text{in}} U(\theta)^\dagger]$  is the ensemble-average trash state, and the measurement operator is  $H_G = \mathbb{1}_{AB} - \mathbb{1}_A \otimes |\mathbf{0}\rangle\langle\mathbf{0}|$ .

Note that  $H_G$  is a global measurement operator, meaning it acts non-trivially on all qubits, which can lead to barren plateaus in the training landscape [8]. To remedy this issue, Ref. [8] proposed a loss function with a local measurement operator acting non-trivially only on a small number of qubits:

$$\mathcal{L}_L(\theta) = 1 - \frac{1}{n_B} \sum_{j=1}^{n_B} \text{Tr}_B[ (|0\rangle\langle 0|_j \otimes \mathbb{1}_{\bar{j}}) \rho_B^{\text{out}} ] \quad (21)$$

$$= \sum_i p_i \text{Tr}_{AB}[H_L U(\theta) |\psi_i\rangle\langle\psi_i| U(\theta)^\dagger], \quad (22)$$

where  $H_L = \mathbb{1}_{AB} - \frac{1}{n_B} \sum_{j=1}^{n_B} \mathbb{1}_A \otimes |0\rangle\langle 0|_j \otimes \mathbb{1}_{\bar{j}}$ , and  $\mathbb{1}_{\bar{j}}$  is the identity on all qubits in  $B$  except the  $j$ -th qubit.

It is clear from (20) and (22) that both loss functions fall under our framework. Namely, they have the form in (12) with  $\ell$  having the linear form in (14).

### 3. Dynamical simulation

Recently a QML-based algorithm was proposed for dynamical simulation [37]. Here the idea is to variationally compile the short-time dynamics into a time-dependent quantum neural network [38]. Then one uses the trained model to extrapolate to longer times. The training data for the compiling process can be taken to be product states, due to a generalization result from Ref. [39].

Let  $U_{\Delta t}$  be a unitary associated with the short-time dynamics. Let  $\{|\Psi_i^P\rangle\}_{i=1}^N$  be a set of product states used for training, where  $|\Psi_i^P\rangle = \bigotimes_{j=1}^n |\psi_{i,j}\rangle$ , and where  $n$  is the number of qubits. Let  $U(\theta)$  be the quantum neural network to be trained. Then the loss function is given by:

$$\mathcal{L}(\theta) = \sum_{i=1}^N \frac{1}{N} \text{Tr}[H_i V(\theta) |\Psi_i^P\rangle\langle\Psi_i^P| V(\theta)^\dagger], \quad (23)$$

where we have defined  $V(\boldsymbol{\theta}) := U(\boldsymbol{\theta})^\dagger U_{\Delta t}$ . Here, the measurement operator is given by  $H_i = \mathbb{1} - \frac{1}{n} \sum_{j=1}^n |\psi_{i,j}\rangle\langle\psi_{i,j}| \otimes \mathbb{1}_{\bar{j}}$ , where  $\bar{j}$  is the set of all qubits excluding the  $j$ -th qubit.

Once again, this loss function clearly falls under our framework, having the form in (12) with  $\ell$  having the linear form in (14).

#### 4. Fidelity for Quantum Neural Networks

Dissipative perceptron-based quantum neural networks (DQNNs) were proposed in Ref. [28] and their trainability was analyzed in Ref. [11]. The loss function is based on the fidelity between the idealized output state and the actual output state of the DQNN. Specifically, we are given access to training data  $\{|\phi_i^{\text{in}}\rangle, |\phi_i^{\text{out}}\rangle\}_{i=1}^N$ , and the DQNN is trained to output a state close to  $|\phi_i^{\text{out}}\rangle$  when the input is  $|\phi_i^{\text{in}}\rangle$ .

For this application, a global loss function was considered with the form

$$\mathcal{L}_G(\boldsymbol{\theta}) = \sum_{i=1}^N \frac{1}{N} \text{Tr}[H_i^G \rho_i^{\text{out}}]. \quad (24)$$

Here,  $\rho_i^{\text{out}} = \mathcal{M}_{\boldsymbol{\theta}}(|\phi_i^{\text{in}}\rangle\langle\phi_i^{\text{in}}|)$  is the output state of the DQNN, which is denoted by  $\mathcal{M}_{\boldsymbol{\theta}}$ . The measurement operator is the projector orthogonal to the ideal output state:  $H_i^G = \mathbb{1} - |\phi_i^{\text{out}}\rangle\langle\phi_i^{\text{out}}|$ .

To avoid the issue of barren plateaus, a local loss function was also considered in Ref. [11]:

$$\mathcal{L}_L(\boldsymbol{\theta}) = \sum_{i=1}^N \frac{1}{N} \text{Tr}[H_i^L \rho_i^{\text{out}}], \quad (25)$$

where the measurement operator is

$$H_i^L = \mathbb{1} - \frac{1}{n_{\text{out}}} \sum_{j=1}^{n_{\text{out}}} |\psi_{i,j}^{\text{out}}\rangle\langle\psi_{i,j}^{\text{out}}| \otimes \mathbb{1}_{\bar{j}}. \quad (26)$$

This loss function is relevant whenever the ideal output states have a tensor-product form across the  $n_{\text{out}}$  output qubits, i.e., of the form  $|\phi_i^{\text{out}}\rangle = |\psi_{i,1}^{\text{out}}\rangle \otimes \dots \otimes |\psi_{i,n_{\text{out}}}^{\text{out}}\rangle$ .

Clearly, these two loss functions fall under our framework, having the form in (12) with  $\ell$  having the linear form in (14).

#### 5. Variational quantum state eigensolver

Near-term methods for quantum principal component analysis have recently been proposed, including the variational quantum state eigensolver (VQSE) [31] and the variational quantum state diagonalization (VQSD) algorithm. Let us first discuss VQSE.

The goals of VQSE are to estimate the  $m$  largest eigenvalues of a density matrix  $\rho$  and to find quantum circuits that prepare the associated eigenvectors. When

combined with a method to prepare the covariance matrix as a density matrix, VQSE can be used for principal component analysis. Note that such a method was proposed in Ref. [40], where it was shown that choosing  $\rho = \sum_i p_i |\psi_i\rangle\langle\psi_i|$  prepares the covariance matrix for a given quantum dataset  $\{p_i, |\psi_i\rangle\}$ .

The VQSE loss function can be written as an energy:

$$\mathcal{L}(\boldsymbol{\theta}) = \text{Tr}[HU(\boldsymbol{\theta})\rho U(\boldsymbol{\theta})^\dagger] \quad (27)$$

$$= \sum_i p_i \text{Tr}[HU(\boldsymbol{\theta})|\psi_i\rangle\langle\psi_i|U(\boldsymbol{\theta})^\dagger] \quad (28)$$

where  $U(\boldsymbol{\theta})$  is a parameterized unitary that is trained to approximately diagonalize  $\rho$ . Note that we inserted the formula  $\rho = \sum_i p_i |\psi_i\rangle\langle\psi_i|$  in order to arrive at (28).

The measurement operator  $H$  is chosen to be non-degenerate over its  $m$ -lowest energy levels. For example, one can choose a global version of this operator:

$$H = \mathbb{1} - \sum_{j=1}^m r_j |e_j\rangle\langle e_j|, \quad r_j > 0 \quad (29)$$

or a local version of this operator:

$$H = \mathbb{1} - \sum_{j=1}^n r_j Z_j, \quad r_j \in \mathbb{R}, \quad (30)$$

where  $Z_j$  is the Pauli- $z$  operator on the  $j$ -th qubit, and with appropriately chosen real coefficients  $r_j$  to achieve non-degeneracy over the  $m$ -lowest energy levels. Regardless, it is clear that the loss function in (28) falls under our framework, having the form in (12) with  $\ell$  having the linear form in (14).

#### 6. Variational quantum state diagonalization

The goal of the VQSD algorithm [30] is essentially the same as that of VQSE, i.e., to diagonalize a target quantum state  $\rho$ . Let us use:

$$\tilde{\rho} := U(\boldsymbol{\theta})\rho U(\boldsymbol{\theta})^\dagger \quad (31)$$

to denote the state after the attempted diagonalization. Here we omit the  $\boldsymbol{\theta}$  dependency for simplicity of notation.

In contrast to VQSE, the VQSD loss function depends quadratically on the quantum state. Specifically, global and local functions have been proposed, respectively given by

$$\mathcal{L}_G(\boldsymbol{\theta}) = \text{Tr}[\rho^2] - \text{Tr}[\mathcal{Z}(\tilde{\rho})^2], \quad (32)$$

$$\mathcal{L}_L(\boldsymbol{\theta}) = \text{Tr}[\rho^2] - \frac{1}{n} \sum_{j=1}^n \text{Tr}[\mathcal{Z}_j(\tilde{\rho})^2]. \quad (33)$$

Here,  $\mathcal{Z}$  and  $\mathcal{Z}_j$  are quantum channels that dephase (i.e., destroy the off-diagonal elements) in the global standard basis and in the local standard basis on qubit  $j$ , respectively.

We can rewrite the terms in the global loss using:

$$\text{Tr}[\rho^2] = \text{Tr}[(\rho \otimes \rho)\text{SWAP}], \quad (34)$$

$$\text{Tr}[\mathcal{Z}(\bar{\rho})^2] = \text{Tr}[(\rho \otimes \rho)W_G^\dagger(|\mathbf{0}\rangle\langle\mathbf{0}| \otimes \mathbf{1})W_G], \quad (35)$$

where *SWAP* denotes the swap operator, and where  $W_G$  corresponds to the layers of CNOTs used in the so-called diagonalized inner product (DIP) test circuit [30]. Hence, we obtain:

$$\mathcal{L}_G(\boldsymbol{\theta}) = \text{Tr}[(\rho \otimes \rho)H_G] \quad (36)$$

$$= \sum_{i,i'} p_i p_{i'} \text{Tr}[(|\psi_i\rangle\langle\psi_i| \otimes |\psi_{i'}\rangle\langle\psi_{i'}|)H_G] \quad (37)$$

where  $H_G = \text{SWAP} - U_G^\dagger(|\mathbf{0}\rangle\langle\mathbf{0}| \otimes \mathbf{1})U_G$ . Note that we inserted the relation  $\rho = \sum_i p_i |\psi_i\rangle\langle\psi_i|$  in order to arrive at (37). Also note that one can think of the  $q_i := p_i p_{i'}$  as defining a probability distribution over the index  $\mathbf{i} = \{i, i'\}$ . Hence it is clear that (37) falls under our framework, with  $m = 2$  copies of the input Hilbert space, and with  $\ell$  having the linear form in (14).

Similarly, for the local loss, we can write

$$\mathcal{L}_L(\boldsymbol{\theta}) = \text{Tr}[(\rho \otimes \rho)H_L] \quad (38)$$

$$= \sum_{i,i'} p_i p_{i'} \text{Tr}[(|\psi_i\rangle\langle\psi_i| \otimes |\psi_{i'}\rangle\langle\psi_{i'}|)H_L] \quad (39)$$

where  $H_L = \text{SWAP} - (1/n) \sum_{j=1}^n H_{L,j}$ , and  $H_{L,j}$  is the Hermitian operator that is measured for the partial diagonalized inner product (PDIP) test [30]. Once again, the local loss falls under our framework, with  $m = 2$  copies of the input Hilbert space, and with  $\ell$  having the linear form in (14).

### 7. Mean squared error for quantum classifiers

The mean squared error (MSE) loss function is widely employed in classical machine learning. Moreover, it has been used in the context of quantum neural networks, e.g., in Ref. [4]. Given a set of labels  $y_i \in \mathbb{R}$  for a dataset  $\{\rho_i\}_{i=1}^N$ , and a set of predictions from a machine learning model denoted  $\tilde{y}_i(\boldsymbol{\theta})$ , the MSE loss is computed as:

$$\mathcal{L}(\boldsymbol{\theta}) = \frac{1}{N} \sum_{i=1}^N (y_i - \tilde{y}_i(\boldsymbol{\theta}))^2. \quad (40)$$

In the case of a quantum model, there is freedom in specifying how to compute the prediction  $\tilde{y}_i(\boldsymbol{\theta})$ . Typically, this will be estimated via an expectation value, such as  $\tilde{y}_i(\boldsymbol{\theta}) = \text{Tr}[\mathcal{M}_{\boldsymbol{\theta}}(\rho_i)H_i] = E_i(\boldsymbol{\theta})$ . In this case, the loss function in (40) would be a quadratic function of expectation value  $E_i(\boldsymbol{\theta})$ . Hence, this falls under our framework, with  $\ell$  having the polynomial form in (15) with degree  $D = 2$ .

### C. Possible extensions to our framework

Up until now, we have considered loss functions that have a linear or polynomial dependence on measurable expectation values,  $E_i(\boldsymbol{\theta})$ . As shown above, this form encompasses many loss functions used in the literature and we will show that cost functions of this form enable cheap gradient estimation and therefore shot frugal optimization.

However, there are also more complicated loss functions that have a less trivial dependence on measurable expectation values. One such loss function is the log-likelihood which is commonly used in classification tasks [5, 41].

Let us assume one is provided with a dataset of the form  $\mathcal{S} = \{\rho_i, y_i\}_{i=1}^N$  where  $\rho_i$  are quantum states, and  $y_i$  are discrete real-valued labels. We aim to train a parametrized model whose goal is to predict the correct labels. Here one cannot directly apply the techniques previously used, as discrete outputs mean we cannot simply evaluate gradients. Therefore we cannot directly employ the results in [24, 35]. Hence, a different shot-frugal method must be employed.

Given  $\mathcal{S}$ , the likelihood function is defined as

$$\text{Pr}(\mathcal{S}|\boldsymbol{\theta}) := \prod_{i=1}^N \text{Pr}(y_i|\rho_i; \boldsymbol{\theta}) = \prod_{i=1}^N p_i, \quad (41)$$

where we have defined  $p_i = \text{Pr}(y_i|\rho_i; \boldsymbol{\theta})$  and where it is assumed that the model's predictions only depend on the current input. From here, we can define the negative log-likelihood loss function as

$$\mathcal{L}(\boldsymbol{\theta}) = -\log(\text{Pr}(\mathcal{S}|\boldsymbol{\theta})) = -\sum_{i=1}^N \log(p_i). \quad (42)$$

In Appendix A we present a conceptually different approach for shot frugal optimization, which does not rely on the unbiased estimation of gradients. Rather, this approach relies on the construction of an unbiased estimator for the log-likelihood loss function. In particular, we follow the results in [42] and use inverse binomial sampling in the context of quantum machine learning. We note that this extension is of independent interest to gradient-free shot-frugal optimizers (as one cannot leverage gradients for problems with discrete outputs), but leave further exploration to future work.

### IV. UNBIASED ESTIMATORS FOR GRADIENTS OF QML LOSSES

Gradient-based optimizers are commonly used when optimizing QML models. In shot frugal versions of these approaches, one of the key aspects is the construction of an unbiased estimator of the gradient [24–26]. There are two types of loss functions we will consider in this work;

those that have a linear dependence on expectation values and those that have a non-linear dependence given by a polynomial function. We will see that for these two types of losses one can construct unbiased estimators of the gradients and that it is also possible to define sampling strategies that massively reduce the number of shots needed to evaluate such an estimator.

Previous work considered how to construct unbiased estimators of QML gradients. However, the shot frugal resource allocation strategies presented were sub-optimal [26] and leave room for improvement. Furthermore, the more sophisticated shot allocation methods presented in [27] were purpose-built for VQE-type cost functions. Here we unify approaches from these two works and show how one can employ more sophisticated shot allocation strategies in a general QML setting.

First, we consider the simpler case of linear loss functions where we show one can directly employ the parameter shift rule to construct an unbiased estimator for the gradient. Then we turn our attention to polynomial loss functions where we present a general form for an unbiased estimator of the gradient. In both cases, we show that shots can be allocated according to the expansion coefficients in the expressions we derive. These in turn depend on the coefficients in the operators to be measured and the set of quantum states used in the QML data set. Such shot-allocation schemes are an important ingredient in the design of our QML shot-frugal optimizer in the next section. We use the U-statistic formalism [43] to construct unbiased estimators for the loss function in each case. For an introduction to U-statistics, we refer the reader to Appendix B.

### A. Loss functions linear in the quantum circuit observables

#### 1. Using the parameter shift rule

Loss functions that have a linear dependence on expectation values of observables are straightforward to consider. As shown in Ref. [26], the parameter shift rule can be used directly to construct unbiased estimators of the gradient of these loss functions. Previously, we wrote our general linear loss function as

$$\mathcal{L}(\boldsymbol{\theta}) = \sum_i p_i \ell(E_i(\boldsymbol{\theta})). \quad (43)$$

Let us not consider the case where

$$E_i(\boldsymbol{\theta}) = \text{Tr}(\mathcal{M}_\theta(\rho_i)H_i) \quad \text{and} \quad \ell(E_i(\boldsymbol{\theta})) = E_i(\boldsymbol{\theta}).$$

By expanding the measurement operator as  $H_i = \sum_{j=1}^{t_i} c_{i,j} h_{i,j}$ , we can then write this loss function as follows

$$\mathcal{L}(\boldsymbol{\theta}) = \sum_{i,j} q_{i,j} \langle h_{i,j}(\boldsymbol{\theta}) \rangle, \quad (44)$$

where  $q_{i,j} = p_i c_{i,j}$  and  $\text{Tr}[\mathcal{M}_\theta(\rho_i)h_{i,j}] = \langle h_{i,j}(\boldsymbol{\theta}) \rangle$ .

Consider the partial derivative of this loss function with respect to the parameter  $\theta_x$ ,

$$\frac{\partial \mathcal{L}}{\partial \theta_x} = \sum_{i,j} q_{i,j} \frac{\partial \langle h_{i,j}(\boldsymbol{\theta}) \rangle}{\partial \theta_x}. \quad (45)$$

Suppose we assume that the quantum channel  $\mathcal{M}_\theta(\rho_i)$  is a unitary channel:

$$\mathcal{M}_\theta(\rho_i) = U(\boldsymbol{\theta})\rho_i U(\boldsymbol{\theta})^\dagger, \quad (46)$$

where  $U(\boldsymbol{\theta})$  is a trainable quantum circuit whose parametrized gates are generated by Pauli operators. This assumption allows us to directly apply the parameter shift rule (see Sec. II). This leads to

$$\frac{\partial \mathcal{L}}{\partial \theta_x} = \frac{1}{2} \sum_{i,j} q_{i,j} (\langle h_{i,j}(\boldsymbol{\theta} + \boldsymbol{\delta}_x \frac{\pi}{2}) \rangle - \langle h_{i,j}(\boldsymbol{\theta} - \boldsymbol{\delta}_x \frac{\pi}{2}) \rangle). \quad (47)$$

Therefore, an unbiased estimator for the gradient can be obtained by combining two unbiased estimators for the loss function. Defining  $\hat{g}_x(\boldsymbol{\theta})$  to be an unbiased estimator of the  $x$ -th component of the gradient,

$$\hat{g}_x(\boldsymbol{\theta}) = \frac{1}{2} [\hat{\mathcal{L}}(\boldsymbol{\theta} + \boldsymbol{\delta}_x \frac{\pi}{2}) - \hat{\mathcal{L}}(\boldsymbol{\theta} - \boldsymbol{\delta}_x \frac{\pi}{2})], \quad (48)$$

where  $\hat{\mathcal{L}}(\boldsymbol{\theta} \pm \boldsymbol{\delta}_x \frac{\pi}{2})$  are unbiased estimators for the loss function at the different shifted parameter values needed when employing the parameter shift rule. This means that for loss functions that are linear in the expectation values recovered from the quantum circuit, one can then use an unbiased estimator for the cost evaluated at different parameter values to return an unbiased estimator for the gradient.

We can therefore distribute the shots according to the coefficients  $q_{i,j}$  when evaluating a single or multi-shot estimate of  $\hat{\mathcal{L}}(\boldsymbol{\theta} \pm \boldsymbol{\delta}_x \frac{\pi}{2})$ . This can be done by constructing a probability distribution according to the probabilities  $\epsilon_{i,j} = |q_{i,j}| / \sum_{i,j} |q_{i,j}|$ . Note that this distribution strategy relies on the construction of two unbiased estimators of the loss function, which are then combined. In the next section, we show how one can construct such estimators. In practice, we will use the same total number of shots to evaluate both estimators as they have the same expansion weights and are therefore equally important.

#### 2. Unbiased estimators for Linear loss functions

Let  $s_{\text{tot}}$  represent the total number of shots. We denote  $\hat{\mathcal{E}}_{i,j}$  as the estimator for  $\langle h_{i,j}(\boldsymbol{\theta}) \rangle$ , and  $\hat{\mathcal{L}}(\boldsymbol{\theta})$  the estimator for the loss. That is,

$$\hat{\mathcal{L}}(\boldsymbol{\theta}) = \sum_{i,j} q_{i,j} \hat{\mathcal{E}}_{i,j}, \quad \text{with} \quad \hat{\mathcal{E}}_{i,j} = \frac{1}{\mathbb{E}[s_{i,j}]} \sum_{k=1}^{s_{i,j}} r_{i,j,k}. \quad (49)$$

Here,  $s_{i,j}$  is the number of shots allocated to the measurement of  $\langle h_{i,j}(\boldsymbol{\theta}) \rangle$ . Note that  $s_{i,j}$  may be a random variable. As we will work in terms of the total shot budget for the estimation,  $s_{\text{tot}}$ , we impose  $\sum_{i,j} s_{i,j} = s_{\text{tot}}$ . Also, each  $r_{i,j,k}$  is an independent random variable associated with the  $k$ -th single-shot measurement of  $\langle h_{i,j}(\boldsymbol{\theta}) \rangle$ . We will assume that  $\mathbb{E}[s_{i,j}] > 0$  for all  $i,j$ . Using these definitions we can show that this defines an unbiased estimator for the loss function in the following proposition.

**Proposition 1.** *Let  $\widehat{\mathcal{L}}$  be the estimator defined in Eq. (49).  $\widehat{\mathcal{L}}$  is an unbiased estimator for the cost function  $\mathcal{L}(\boldsymbol{\theta})$  defined in Eq. (12).*

*Proof.*

$$\begin{aligned} \mathbb{E}[\widehat{\mathcal{L}}] &= \mathbb{E}\left[\sum_{i,j} q_{i,j} \frac{1}{\mathbb{E}[s_{i,j}]} \sum_{k=1}^{s_{i,j}} r_{i,j,k}\right] \\ &= \sum_{i,j} q_{i,j} \frac{1}{\mathbb{E}[s_{i,j}]} \mathbb{E}\left[\sum_{k=1}^{s_{i,j}} r_{i,j,k}\right]. \end{aligned} \quad (50)$$

Here it is useful to recall Wald's equation. Wald's equation states that the expectation value of the sum of  $N$  real-valued, identically distributed, random variables,  $X_i$ , can be expressed as

$$\mathbb{E}\left[\sum_{i=1}^N X_i\right] = \mathbb{E}[N]\mathbb{E}[X_1], \quad (51)$$

where  $N$  is a random variable that does not depend on the terms of the sum. In our case, each shot is indeed independent and sampled from the same distribution. Furthermore, the total number of shots does not depend on the sequence of single-shot measurements. Therefore,

$$\mathbb{E}[\widehat{\mathcal{L}}] = \sum_{i,j} q_{i,j} \frac{\mathbb{E}[s_{i,j}]}{\mathbb{E}[s_{i,j}]} \langle h_{i,j}(\boldsymbol{\theta}) \rangle = \mathcal{L}(\boldsymbol{\theta}). \quad (52)$$

□

Alternatively, one can see that the estimator defined in Eq. (49) has the form of a degree-1 U-statistic for  $\widehat{\mathcal{L}}$ . We the above result, we arrive at the following corollary.

**Corollary 1.** *Let  $\widehat{g}_x(\boldsymbol{\theta})$  be the estimator defined in Eq. (48).  $\widehat{g}_x(\boldsymbol{\theta})$  is an unbiased estimator for the  $x$ -th component of the gradient.*

*Proof.* The proof follows by taking the expectation values of Eq. (48) and employing the result from Prop. 1 and the parameter shift rule. □

We also obtained the variance of the estimator constructed above whose derivation can be found in Appendix C. Such a quantity can be used to compare different strategies for allocating shots among Hamiltonian terms. Although we do not compare variances, we use elements of the proof in Appendix C for the following

construction of an unbiased estimator of the MSE in Appendix E.

Having now constructed an estimator for the loss function as outlined in the previous section combining two single or multi-shot estimates of this function evaluated at the required parameter values will lead to an unbiased estimator of the gradient.

## B. Loss functions with polynomial dependence on the quantum circuit observables

### 1. Constructing an unbiased estimator of the gradient

In the case of non-linear dependence on the observables produced by a quantum circuit, estimating the gradient is not as simple as simply applying to parameter shift rule. However, we can still derive estimators for the gradient when the non-linearity is described by a polynomial function. We begin with the general expression for the loss functions we consider in this work

$$\mathcal{L}(\boldsymbol{\theta}) = \sum_i p_i \ell(E_i(\boldsymbol{\theta})). \quad (53)$$

Now constructing a polynomial loss function of degree  $D$  requires that

$$\ell(E_i(\boldsymbol{\theta})) = \sum_{z=0}^D a_z [E_i(\boldsymbol{\theta})]^z, \quad (54)$$

which leads to the expression of the loss function,

$$\mathcal{L}(\boldsymbol{\theta}) = \sum_{i,z} p_{i,z} \left[ \sum_j c_{i,j} \langle h_{i,j}(\boldsymbol{\theta}) \rangle \right]^z, \quad (55)$$

where  $p_{i,z} = p_i a_z$ . Taking the derivative with respect to  $\theta_x$  leads to

$$\frac{\partial \mathcal{L}(\boldsymbol{\theta})}{\partial \theta_x} = \sum_{i,z} p_{i,z} z \left( \sum_j c_{i,j} \langle h_{i,j}(\boldsymbol{\theta}) \rangle \right)^{z-1} \left( \sum_{j'} c_{i,j'} \frac{\partial \langle h_{i,j'}(\boldsymbol{\theta}) \rangle}{\partial \theta_x} \right). \quad (56)$$

Using the multinomial theorem and given  $J$  Hamiltonian terms, we can expand the second sum in the above expression,

$$\begin{aligned} \frac{\partial \mathcal{L}(\boldsymbol{\theta})}{\partial \theta_x} &= \sum_{i,z} p_{i,z} z \sum_{b_1+b_2+\dots+b_J=z-1} \binom{z-1}{b_1, b_2, \dots, b_J} \\ &\quad \prod_j (c_{i,j} \langle h_{i,j}(\boldsymbol{\theta}) \rangle)^{b_j} \left( \sum_{j'} c_{i,j'} \frac{\partial \langle h_{i,j'}(\boldsymbol{\theta}) \rangle}{\partial \theta_x} \right), \end{aligned} \quad (57)$$

where  $\binom{z-1}{b_1, b_2, \dots, b_J} = \frac{(z-1)!}{b_1! b_2! \dots b_J!}$  and  $b_j$  are non-negative integers. Therefore, we need to construct unbiased estimators of the terms  $\langle h_{i,j}(\boldsymbol{\theta}) \rangle^{b_j}$  and use the previously established gradient estimators with the parameter shift

rule. Then we will need to consider how to distribute shots among this estimator. Rewriting Eq. (57) leads to

$$\frac{\partial \mathcal{L}(\boldsymbol{\theta})}{\partial \theta_x} = \sum_{\mathbf{i}, z} \sum_{b_1 + b_2 + \dots + b_J = z-1} p_{\mathbf{i}, z} z \frac{(z-1)!}{b_1! b_2! \dots b_J!} \sum_{j'} \prod_j c_{\mathbf{i}, j}^{b_j} c_{\mathbf{i}, j'} \left[ \frac{\partial \langle h_{\mathbf{i}, j'}(\boldsymbol{\theta}) \rangle}{\partial \theta_x} \langle h_{\mathbf{i}, j}(\boldsymbol{\theta}) \rangle^{b_j} \right]. \quad (58)$$

Therefore we can distribute the shots according to the magnitude of the expansion terms  $\prod_j c_{\mathbf{i}, j}^{b_j} c_{\mathbf{i}, j'} p_{\mathbf{i}, z} z \frac{(z-1)!}{b_1! b_2! \dots b_J!}$ . Once again we can construct an unbiased estimator with the normalized magnitude of these terms defining the probabilities. We now explore how to construct an unbiased estimator for the term  $\langle h_{\mathbf{i}, j}(\boldsymbol{\theta}) \rangle^{b_j}$ , which is essential to construct an unbiased estimator for the gradients of these kinds of loss functions.

## 2. Constructing an unbiased estimator of polynomial terms

In order to construct an unbiased estimator for the gradient we need an unbiased estimator for terms of the form

$$\frac{\partial \langle h_{\mathbf{i}, j}(\boldsymbol{\theta}) \rangle}{\partial \theta_x} \prod_j \langle h_{\mathbf{i}, j}(\boldsymbol{\theta}) \rangle^{b_j}. \quad (59)$$

We can use the parameter shift rule for the term on the left-hand side. For the term in the product as each  $j$  index corresponds to a different operator in the Hamiltonian, each term will be independent. Therefore, we need to construct estimators for terms of the form  $\langle h_{\mathbf{i}, j}(\boldsymbol{\theta}) \rangle^z$ . This leads us to the following proposition.

**Proposition 2.** Let  $\widehat{\xi}_{\mathbf{i}, j}$  be an estimator defined as

$$\widehat{\xi}_{\mathbf{i}, j} = \frac{1}{\mathbb{E}[\binom{s_{\mathbf{i}, j}}{z}]} \sum h^*(r_{\mathbf{i}, j, k_{\alpha_1}}, \dots, r_{\mathbf{i}, j, k_{\alpha_z}}), \quad (60)$$

where the summation is over all subscripts  $1 \leq \alpha_1 < \alpha_2 < \dots < \alpha_z \leq s_{\mathbf{i}, j}$  and  $h^*(r_{\mathbf{i}, j, k_1}, \dots, r_{\mathbf{i}, j, k_z}) = \prod_{\beta=1}^z r_{\mathbf{i}, j, k_{\alpha_\beta}}$ .  $\widehat{\xi}_{\mathbf{i}, j}$  is an unbiased estimator for the term  $\langle h_{\mathbf{i}, j}(\boldsymbol{\theta}) \rangle^z$  estimated with  $s_{\mathbf{i}, j}$  shots where  $s_{\mathbf{i}, j} \geq z$ .

*Proof.* Eq. (60) is inspired by the form of a U-statistic for  $\langle h_{\mathbf{i}, j}(\boldsymbol{\theta}) \rangle^{b_j}$ , see Appendix B for more details. Taking the expectation value and using the U-statistic formalism one arrives at the desired result.  $\square$

Bringing this all together we can formulate a proposition regarding an unbiased estimator for gradients of loss functions with polynomial dependence.

**Proposition 3.**

$$\hat{g}_x(\boldsymbol{\theta}) = \sum_{\mathbf{i}, z} \sum_{b_1 + b_2 + \dots + b_J = z-1} p_{\mathbf{i}, z} z \frac{(z-1)!}{b_1! b_2! \dots b_J!} \sum_{j'} \prod_{j=1}^J c_{\mathbf{i}, j}^{b_j} c_{\mathbf{i}, j'} \frac{1}{\mathbb{E}[s_{\mathbf{i}, j, j', \mathbf{b}}]} \sum_{k=1}^{s_{\mathbf{i}, j, j', \mathbf{b}}} (r_{\mathbf{i}, j, k}^+ - r_{\mathbf{i}, j, k}^-) \frac{1}{\mathbb{E}[\binom{s_{\mathbf{i}, j, j', \mathbf{b}}}{b_j}]} \sum_{\beta=1}^z r_{\mathbf{i}, j, k_{\alpha_\beta}}, \quad (61)$$

where

$$\mathbb{E}[r_{\mathbf{i}, j, 1}^\pm] = \langle h_{\mathbf{i}, j}(\boldsymbol{\theta} \pm \boldsymbol{\delta}_x \frac{\pi}{2}) \rangle \quad (62)$$

and  $\mathbf{b} = (b_1, b_2, \dots, b_J)$ . The final sum is over all subscripts  $1 \leq \alpha_1 < \alpha_2 < \dots < \alpha_z \leq s_{\mathbf{i}, j}$  and  $h^*(r_{\mathbf{i}, j, k_1}, \dots, r_{\mathbf{i}, j, k_z}) = \prod_{\beta=1}^z r_{\mathbf{i}, j, k_{\alpha_\beta}}$ .  $\hat{g}_x(\boldsymbol{\theta})$  is a unbiased estimator for  $\frac{\partial \mathcal{L}(\boldsymbol{\theta})}{\partial \theta_x}$

*Proof.* This can be seen to be true by taking the expectation values and invoking the propositions previously established.  $\square$

These same techniques can be used to construct unbiased estimators of the loss function, which we expand on in Appendix D. In order to provide a concrete example of using the above propositions, we consider the special case of the MSE loss function.

## 3. Constructing an estimator for the gradient of the MSE loss function

To clarify the notation used above, let us focus on the special case of the MSE loss function. We consider a slightly more general form than the MSE cost introduced above in Eq. (40). Consider a set of labels  $y_i \in \mathbb{R}$  for a dataset  $\{\rho_i\}_{i=1}^N$ , composed of the tensor product of  $m$  states. The set of predictions from the quantum machine learning model are denoted  $\tilde{y}_i(\boldsymbol{\theta}) = \sum_j c_{\mathbf{i}, j} \langle h_{\mathbf{i}, j}(\boldsymbol{\theta}) \rangle$ . Therefore, we can write the loss as

$$\mathcal{L}_{MSE}(\boldsymbol{\theta}) = \sum_i p_i [y_i - \sum_j c_{\mathbf{i}, j} \langle h_{\mathbf{i}, j}(\boldsymbol{\theta}) \rangle]^2. \quad (63)$$

Expanding the previous equation leads to

$$\mathcal{L}_{MSE}(\boldsymbol{\theta}) = \sum_i p_i [y_i^2 - y_i \sum_j c_{\mathbf{i}, j} \langle h_{\mathbf{i}, j}(\boldsymbol{\theta}) \rangle + (\sum_j c_{\mathbf{i}, j} \langle h_{\mathbf{i}, j}(\boldsymbol{\theta}) \rangle)^2]. \quad (64)$$

Evaluating the partial derivative with respect to  $\theta_x$  gives,

$$\frac{\partial \mathcal{L}_{MSE}(\boldsymbol{\theta})}{\partial \theta_x} = \sum_{\mathbf{i}, j, j'} -p_i c_{\mathbf{i}, j} \frac{\partial \langle h_{\mathbf{i}, j}(\boldsymbol{\theta}) \rangle}{\partial \theta_x} + 2p_i c_{\mathbf{i}, j'} c_{\mathbf{i}, j} \langle h_{\mathbf{i}, j'}(\boldsymbol{\theta}) \rangle \frac{\partial \langle h_{\mathbf{i}, j}(\boldsymbol{\theta}) \rangle}{\partial \theta_x}. \quad (65)$$

We can distribute the total number of shots  $s_{tot}$  among these terms according to the relative magnitudes of the  $p_i c_{i,j'}$  and  $2p_i c_{i,j'} c_{i,j}$  coefficients. This can be achieved by sampling from a multinomial probability distribution where the normalized magnitude of these coefficients defines the probabilities.

It is important to note that some estimators have a different minimum number of shots than others. For example, the estimator for the last term in the above equation, involving a gradient and a direct expectation value, requires a total of 3 for different circuit evaluations. The estimator for this term can be written as

$$\widehat{\mathcal{D}}_{i,j,j'} = \frac{1}{2} \widehat{\mathcal{E}}_{i,j'} (\widehat{\mathcal{E}}_{i,j}^+ - \widehat{\mathcal{E}}_{i,j}^-). \quad (66)$$

Therefore, in this case, the minimum number of shots given to any term can be set to 3 to ensure every estimation of any term in Eq. (63) will always be unbiased. We expand on this consideration below.

### C. Distributing the shots among estimator terms

As previously mentioned the shots can be distributed according to a multinomial distribution with probabilities given by the magnitude of the constant factors that appear in the expression for the above estimators. We note that the shots assigned to a given term may be zero. In order to ensure the estimate of the above term using a given number of shots we need to ensure the estimate of each term is also unbiased. One needs at least 2 shots to produce an unbiased estimate of the gradient. For terms of the form  $\langle h_{i,j}(\boldsymbol{\theta}) \rangle^{b_j}$  one needs at least  $b_j$  shots, corresponding to the degree of U-statistic. Therefore, care needs to be taken when distributing shots to ensure that each term measured has sufficiently many. One way to ensure this is the case is to distribute shots in multiples of the largest number of shots needed to evaluate any one term. Any leftover shots can then be distributed equally across the terms to be measured.

Each term itself may consist of a product of several expectation values, each with a different required minimum number of shots. The shots distribution with each term can be selected to correspond to this required minimum number of shots per term. To make this concrete consider a term of the form in Eq. (59). Estimating the gradient will take at least 2 shots. Estimating the product term will take at least  $\sum_{j=1}^J b_j$  shots. If we are given  $s_{i,j}$  shots to use to estimate this term we can assign  $\lfloor 2(s_{i,j}/(2 + \sum_{j=1}^J b_j)) \rfloor$  to the first term and  $\lfloor \sum_{j=1}^J b_j (s_{i,j}/(2 + \sum_{j=1}^J b_j)) \rfloor$  to the product term, distributing any remaining shots equally among both.

## V. THE REFOQUS OPTIMIZER

Now that we have defined unbiased estimators of the gradient, we have the tools needed to present our new

optimizer. We call our optimizer Refoqus, which stands for REsource Frugal Optimizer for QUantum Stochastic gradient descent. This optimizer is tailored to QML tasks. QML tasks have the potential to incur large shot overheads because of large training datasets as well as measurement operators composed of a large number of terms. With Refoqus, we achieve shot frugality first by leveraging the gCANS [25] rule for allocating shots at each iteration step and second by allocating these shots to individual terms in the loss function via random sampling over the training dataset and over measurement operators. This random sampling allows us to remove the shot floor imposed by deterministic strategies (while still achieving an unbiased estimator for the gradient), hence unlocking the full shot frugality of our optimizer.

The key insight needed for the construction of the Refoqus protocol is that cost functions that have a linear or polynomial dependence on measurable hermitian matrices, and their gradients can always be written in a form consisting of a summation of different measurable quantities with expansion coefficients. These coefficients can in turn be used to define a multinomial probability distribution to guide the allocation of the shots to each term when evaluating the loss function or its gradient.

We outline the Refoqus optimizer in Algorithm 1. Given a number of shots to distribute among Hamiltonian terms, one evaluates the gradient components and the corresponding variances with the *iEvaluate* subroutine (Line 3). We follow the gCANS procedure to compute the shot budget for each iteration (Line 12). We refer to [25] for more details on gCANS. The iterative process stops until the total shot budget has been used.

The hyperparameters that we employ are similar to those of Rosalin [27], and will come with similar recommendations on how to set them. For instance, the Lipschitz constant  $L$  bounds the largest possible value of the derivative. Hence, it depends on the loss expression but can be set as  $M = \sum_{i,j} |q_{i,j}|$  according to [24]. Moreover, a learning rate satisfying  $0 < \alpha < 2/L$  can be used.

## VI. CONVERGENCE GUARANTEES

The framework for Refoqus leverages the structure and the update rule of the gCANS optimizer presented in [25]. Therefore, we can apply the same arguments introduced to show geometric convergence. We repeat the arguments and assumptions needed for this convergence result here for convenience.

**Proposition 4.** *Provided the loss function satisfies the assumptions stated below the stochastic gradient descent routine used in Refoqus achieves geometric convergence to the optimal value of the cost function. That is,*

$$\mathbb{E}[\mathcal{L}(\boldsymbol{\theta})^{(t)}] - \mathcal{L}^* = \mathcal{O}(\gamma^t) \quad (67)$$

where  $t$  labels the iteration,  $\mathcal{L}^*$  is the optimal value of the loss function, and  $0 < \gamma < 1$ .

**Algorithm 1** The optimization loop used in *Refoqus*. The function  $iEvaluate(\boldsymbol{\theta}, \mathbf{s}, \{f(p_i, c_{i,j})\})$  evaluates the the gradient at  $\boldsymbol{\theta}$  returning, a vector of the gradient estimates  $\mathbf{g}$  and their variances  $\mathbf{S}$ . The vectors are calculated using  $s_0 s_x$  shots to estimate the  $x$ -th component of the gradient and its variance, where  $s_0$  is the minimum number of shots required to obtain an unbiased estimator. Note that shots for each component estimation are distributed in multiples of  $s_0$ , according to a multinomial distribution determined by the expansion coefficients that define the gradient estimator. These expansion coefficients are defined by the function  $f(p_i, c_{i,j})$ , which returns the probability of measuring the term corresponding to the coefficients  $p_i, c_{i,j}$ . For the case of linear loss functions  $f(p_i, c_{i,j}) = \frac{|p_i c_{i,j}|}{\sum_{i,j} p_i c_{i,j}}$ .

**Input:** Learning rate  $\alpha$ , starting point  $\boldsymbol{\theta}_0$ , min number of shots per estimation  $s_{\min}$ , number of shots that can be used in total  $s_{\max}$ , Lipschitz constant  $L$ , running average constant  $\mu$ , a vector of the least number of shots needed for each gradient estimate  $s_0$ , which is loss function dependent and  $f(p_i, c_{i,j})$ , which is also loss function dependent.

```

1: initialize:  $\boldsymbol{\theta} \leftarrow \boldsymbol{\theta}_0, s_{\text{tot}} \leftarrow 0, \mathbf{g} \leftarrow (0, \dots, 0)^T, \mathbf{S} \leftarrow (0, \dots, 0)^T, \mathbf{s} \leftarrow (s_{\min}, \dots, s_{\min})^T, \boldsymbol{\chi}' \leftarrow (0, \dots, 0)^T, \boldsymbol{\chi} \leftarrow (0, \dots, 0)^T, \boldsymbol{\xi} \leftarrow (0, \dots, 0)^T, \boldsymbol{\xi}' \leftarrow (0, \dots, 0)^T, t \leftarrow 0$ 
2: while  $s_{\text{tot}} < s_{\max}$  do
3:    $\mathbf{g}, \mathbf{S} \leftarrow iEvaluate(\boldsymbol{\theta}, \mathbf{s}, \{f(p_i, c_{i,j})\})$ 
4:    $s_{\text{tot}} \leftarrow s_{\text{tot}} + \sum_x s_0 s_x$ 
5:    $\boldsymbol{\chi}' \leftarrow \mu \boldsymbol{\chi} + (1 - \mu) \mathbf{g}$ 
6:    $\boldsymbol{\xi}' \leftarrow \mu \boldsymbol{\xi} + (1 - \mu) \mathbf{S}$ 
7:    $\boldsymbol{\xi} \leftarrow \boldsymbol{\xi}' / (1 - \mu^{t+1})$ 
8:    $\boldsymbol{\chi} \leftarrow \boldsymbol{\chi}' / (1 - \mu^{t+1})$ 
9:    $\boldsymbol{\theta} \leftarrow \boldsymbol{\theta} - \alpha \mathbf{g}$ 
10:  for  $x \in [1, \dots, d]$  do
11:     $s_x \leftarrow \left[ \frac{2L\alpha}{2-L\alpha} \frac{\xi_x \sum_x \xi_x}{\|\boldsymbol{\chi}\|^2} \right]$ 
12:  end for
13:   $t \leftarrow t + 1$ 
14: end while

```

The update rule used in this work and introduced in [25] and the proof presented here can be directly applied. This update rule guarantees fast convergence in expectation to the optimal value of sufficiently smooth loss functions. The underlying assumption of this result is that the loss function is strongly convex and has Lipschitz-continuous gradients. In most realistic QML scenarios the optimization landscape is not convex, however, if the optimizer were to settle into a convex region then fast convergence is expected.

The exact assumptions needed in order to ensure the geometric convergence of *Refoqus* to the global minima are as follows:

1.  $\mathbb{E}[g_x(\boldsymbol{\theta})] = \frac{\partial \mathcal{L}(\boldsymbol{\theta})}{\partial \theta_x}, \forall x \in [d]$ .
2.  $\text{Var}[g_x(\boldsymbol{\theta})] = \frac{\text{Var}[X_x]}{s_x}$ , where  $X_x$  is the sampling-based estimator of  $g_x(\boldsymbol{\theta})$ .
3.  $\mathcal{L}(\boldsymbol{\theta})$  is  $\mu$ -strongly convex.




FIG. 4. **Hardware-efficient ansatz with 2 layers used for VQSE on 4 qubits.** Each  $R_y$  rotation is independently parametrized according to  $R_y(\theta) = e^{-iY\theta/2}$ .

4.  $\mathcal{L}(\boldsymbol{\theta})$  has  $L$ -Lipschitz continuous gradient.
5.  $\alpha$  is a constant learning rate satisfying  $0 < \alpha < \min\{1/L, 2/\mu\}$ .
6. An ideal version of the update rule holds, that is:

$$s_x = \frac{2L\alpha}{2-L\alpha} \frac{\sigma_x (\sum_{x'=1}^d \sigma_{x'})}{\|\nabla \mathcal{L}(\boldsymbol{\theta})\|^2} \quad \forall x \in [d],$$

$$\text{where } \sigma_x = \sqrt{\text{Var}[X_x]}.$$

In the previous sections, we have shown how to construct unbiased estimators for the gradient, satisfying the first assumption. The second assumption is satisfied as the estimate of the gradient is constructed by calculating the mean of several sampling-based estimates. Assumption 5 is satisfied by construction. Assumption 6 is an idealized version of the update rule used in *Refoqus*. However, the gradient magnitude and estimator variances cannot be exactly known in general, so these quantities are replaced by exponentially moving averages to predict the values of  $\sigma_x$  and  $\|\nabla \mathcal{L}(\boldsymbol{\theta})\|^2$ . Assumption 4 is also satisfied for all the loss functions we consider in this work. Finally, as previously noted assumption 5 is not expected to hold in QML landscapes, which are non-convex in general. Nevertheless, the convergence guarantee provides strong analytical motivation for the functionality of the optimizer in an idealized scenario.

## VII. NUMERICAL RESULTS

We benchmark the performance of several optimizers when applied to using VQSE [31] for quantum PCA [44] on molecular ground states. We perform quantum PCA on 3 quantum datasets: ground states of the  $H_2$  molecule in the sto-3g basis (4 qubits),  $H_2$  in the 6-31g basis (8 qubits) and  $\text{BeH}_2$  in the sto-3g basis (14 qubits). We use 101 circuits, per dataset. The corresponding covariance matrix [40], which can be expressed as  $\frac{1}{101} \sum_{i=0}^{100} |\psi\rangle_i \langle \psi|_i$ . We optimize using the local cost introduced in Eq. (30) and repeated here for convenience:

$$H = \mathbb{1} - \sum_{j=1}^n r_j Z_j, \quad r_j \in \mathbb{R} \quad (68)$$

taking coefficients  $r_j = 1.0 + 0.2(j - 1)$  and  $p_i = \frac{1}{101}$  following the presentation in [31]. Hence the latter coefficients form the  $q_{i,j} = p_i r_j$  terms used for shot-allocation strategies, as outlined below.

When implementing Rosalin and Refoqus we use the weighted random sampling strategy for allocating shots as it demonstrated superior performance against other strategies in the numerical results of [27]. Hence, the  $q_{i,j}$  coefficients that appear in the loss function (and the gradient estimators) are used to construct a multinomial probability distribution defined by the terms  $\frac{|q_{i,j}|}{M}$ , where  $M = \sum_{i,j} |q_{i,j}|$ . This distribution is used to probabilistically allocate the number of shots given to each term for each gradient estimation.

We use a circuit consisting of 2 layers of a hardware-efficient ansatz with depth 4, depicted in Fig. 4, with 20, 40, and 70 parameters for the 4, 8 and 16 qubit problems respectively. These parameters are optimized in order to minimize the VQSE cost function using the Adam, Rosalin, and Refoqus optimizers<sup>1</sup>. We use the absolute error in estimating the eigenvalues of the system to benchmark the overall performance of the output of the optimized circuit as this is desired output of the algorithm. Given the exact 16 highest eigenvalues  $\lambda_i$ , and their estimation  $\tilde{\lambda}_i$  given current parameters  $\theta$ , the latter is computed as  $\epsilon_\lambda = \sum_{i=1}^{16} (\lambda_i - \tilde{\lambda}_i)^2$ .

Figure 5 shows the results obtained when running each algorithm 20 times with different random initialization of the variational ansatz, up to a total shot budget of  $10^8$ . In general, we see that Refoqus is the best-performing optimizer, achieving a best-case accuracy while using fewer shots for all shot budgets. In summary, for the 4 qubit system, a median eigenvalue error of  $6.8 \times 10^{-6}$ ,  $5.15 \times 10^{-6}$  and  $6.08 \times 10^{-7}$  was obtained using Adam, Rosalin and Refoqus respectively. Although we found better minimal value with Adam compared to Rosalin ( $3.79 \times 10^{-7}$  and  $1.74 \times 10^{-6}$  respectively), Rosalin is more advantageous at lower shot budgets. However, Refoqus achieved a minimal error value of  $6.13 \times 10^{-9}$  and demonstrates a clear advantage over both Adam and Rosalin. On both the 8 and 14 qubit systems, we observe a similar trend although eigenvalue errors are worse overall. This is due to the variational ansatz being kept at a fixed depth while increasing the size of the problem, which leads to worse performance. Nonetheless, Refoqus appears to clearly match or outperform both Adam and Rosalin in the number of shots required to reach a given accuracy.

Additionally, Refoqus requires fewer parameter updates (iterations) when compared to Rosalin. Indeed, for the 4, 8 and 14 qubits problems respectively, a median of 117, 89, and 80 iterations were used by Refoqus

against 1217, 618, and 353 for Rosalin. We note that this may be interpreted as another desirable feature of the optimizer, as this minimizes the number of iterations needed in the quantum-classical feedback loop to arrive at a solution of the same (or better) quality. Although the number of iterations is not a bottleneck in and of itself, in current hardware the quantum-classical feedback can prove restrictive meaning fewer iterations are favorable for real-device implementation.

Finally, we note that the variance of the Refoqus results appears to be larger, suggesting that sampling over more terms can lead to a larger variety in the quality of the optimization obtained. Nevertheless, the advantage in shot frugality that arises from sampling over more terms is clear and despite this larger variance, Refoqus is clearly the best-performing optimizer.

## VIII. DISCUSSION

QML algorithms present a different paradigm for data processing which is particularly well suited to quantum data. VQAs are a key contender in giving a near-term useful quantum advantage. However, both VQAs and QML models require long run times and large resource overheads during training (as many iterations and shots to achieve respectable performances). To address these challenges, we propose Refoqus as a shot-frugal gradient-based optimizer based on state and operator sampling. We outline many cost functions that are of interest to the community and are easily captured within our framework for Refoqus. Our new optimizer leverages the loss function of the problem to allocate shots randomly when evaluating gradients. This randomness allows us to make resource-cheap gradient update steps, unlocking shot-frugality. We have shown that Refoqus comes with geometric convergence guarantees under specific assumptions, elaborated in Section VI. Additionally, when applying our optimizer to a QML task, namely a quantum PCA task, we obtain significantly better performance in terms of the number of shots and the number of iterations needed to obtain a given accuracy.

A potential future research direction is to extend our analysis to more complicated, non-linear loss functions such as the log-likelihood, exploring in more detail how to introduce shot frugality to gradient-free optimizers. Furthermore, applying our optimizer to a problem of interest on a real device is an interesting and logical next step.

## ACKNOWLEDGMENTS

This work was supported by the U.S. Department of Energy (DOE) through a quantum computing program sponsored by the Los Alamos National Laboratory (LANL) Information Science & Technology Institute. M.H.G. was partially supported by the Labora-

<sup>1</sup> For Adam, we perform 100 shots per circuit in our simulations. For Rosalin and Refoqus, we use a minimum of 2 shots per circuit (this leads to  $s_{\min} = 202$  and  $s_{\min} = 2$  for Rosalin and Refoqus respectively) in conjunction with WRS.

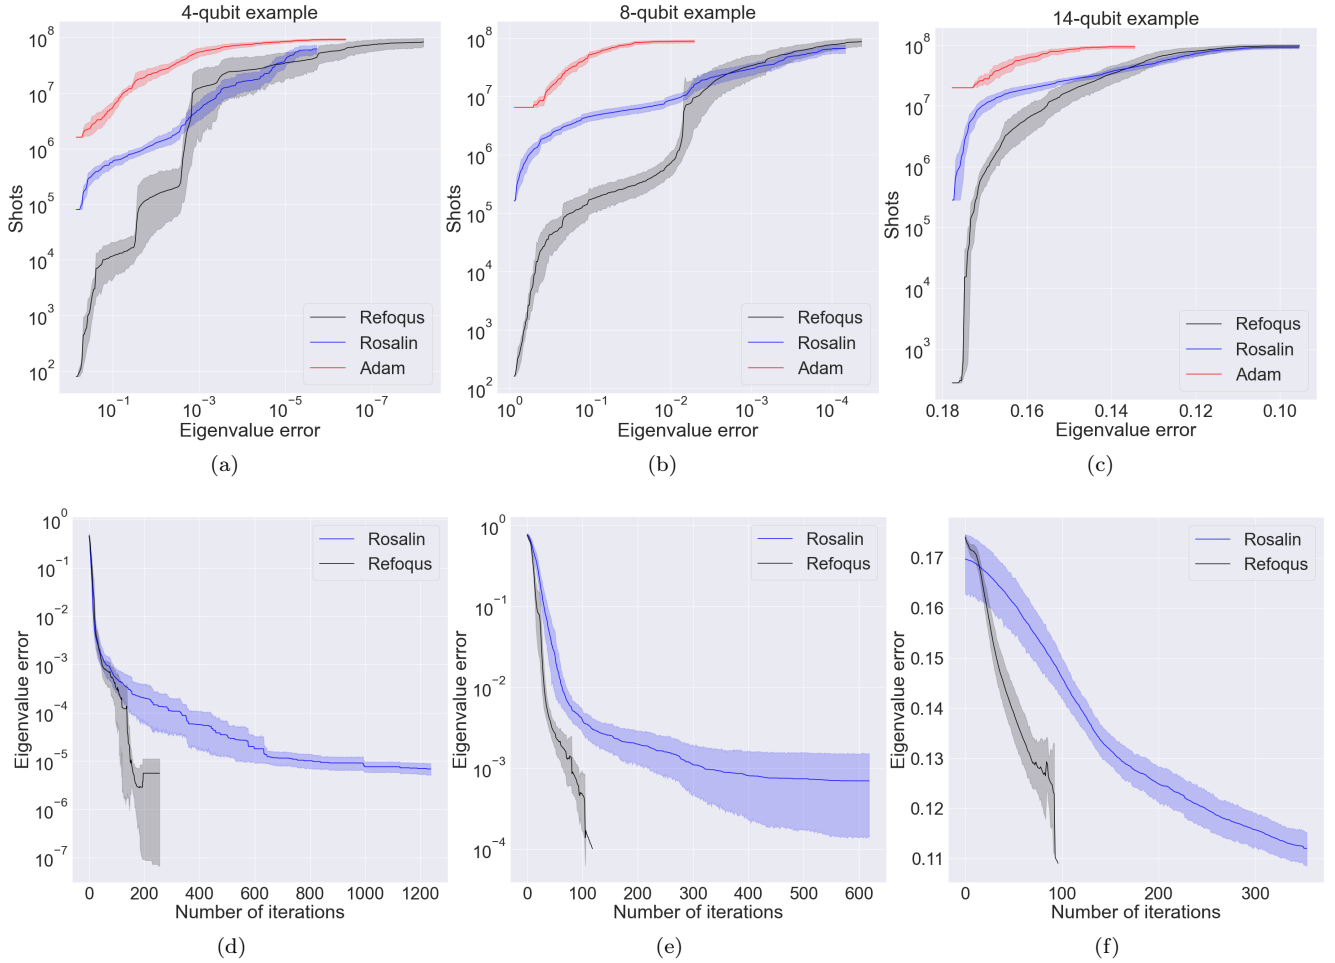


FIG. 5. VQSE for quantum PCA of  $H_2$  molecular ground states (a, b, d, and e) and  $BeH_2$  (c and f). The upper plots show the budget of shots spent against the best eigenvalue error achieved by the optimizers. The lower plots show the number of iterations used to spend the total shot budget. We display the results obtained from 20 independent optimization runs on a data set of 101 circuits representing molecular ground states calculated using the Adam, Rosalin, and Refoqus optimizers and compare their performance. We show the median (solid lines) and 95% confidence intervals (shaded regions) over the 20 different random initializations.

tory Directed Research and Development (LDRD) program under project number 20210116DR. M.C. acknowledges support from the LDRD program of LANL under project number 20230049DR. L.C. was partially supported by the U.S. DOE, Office of Science, Office of Ad-

vanced Scientific Computing Research, under the Accelerated Research in Quantum Computing (ARQC) program. P.J.C. acknowledges support from the LANL ASC Beyond Moore’s Law project.

- 
- [1] J. Biamonte, P. Wittek, N. Pancotti, P. Rebentrost, N. Wiebe, and S. Lloyd, *Nature* **549**, 195 (2017).
- [2] M. Schuld, I. Sinayskiy, and F. Petruccione, *Contemporary Physics* **56**, 172 (2015).
- [3] M. Schuld, I. Sinayskiy, and F. Petruccione, *Quantum Information Processing* **13**, 2567 (2014).
- [4] I. Cong, S. Choi, and M. D. Lukin, *Nature Physics* **15**, 1273 (2019).
- [5] A. Abbas, D. Sutter, C. Zoufal, A. Lucchi, A. Figalli, and S. Woerner, *Nature Computational Science* **1**, 403 (2021).
- [6] Q. T. Nguyen, L. Schatzki, P. Braccia, M. Ragone, M. Larocca, F. Sauvage, P. J. Coles, and M. Cerezo, *arXiv preprint arXiv:2210.08566* (2022).
- [7] J. R. McClean, S. Boixo, V. N. Smelyanskiy, R. Babbush, and H. Neven, *Nature Communications* **9**, 1 (2018).
- [8] M. Cerezo, A. Sone, T. Volkoff, L. Cincio, and P. J. Coles, *Nature Communications* **12**, 1 (2021).
- [9] Z. Holmes, A. Arrasmith, B. Yan, P. J. Coles, A. Albrecht, and A. T. Sornborger, *Physical Review Letters*

**126**, 190501 (2021).

[10] Z. Holmes, K. Sharma, M. Cerezo, and P. J. Coles, *PRX Quantum* **3**, 010313 (2022).

[11] K. Sharma, M. Cerezo, L. Cincio, and P. J. Coles, *Physical Review Letters* **128**, 180505 (2022).

[12] C. O. Marrero, M. Kieferová, and N. Wiebe, *PRX Quantum* **2**, 040316 (2021).

[13] A. Uvarov and J. D. Biamonte, *Journal of Physics A: Mathematical and Theoretical* **54**, 245301 (2021).

[14] A. Arrasmith, M. Cerezo, P. Czarnik, L. Cincio, and P. J. Coles, *Quantum* **5**, 558 (2021).

[15] A. Pesah, M. Cerezo, S. Wang, T. Volkoff, A. T. Sornborger, and P. J. Coles, *Physical Review X* **11**, 041011 (2021).

[16] L. Bittel and M. Kliesch, *Phys. Rev. Lett.* **127**, 120502 (2021).

[17] E. R. Anschuetz and B. T. Kiani, *arXiv preprint arXiv:2205.05786* (2022).

[18] S. Wang, E. Fontana, M. Cerezo, K. Sharma, A. Sone, L. Cincio, and P. J. Coles, *Nature Communications* **12**, 1 (2021).

[19] D. Stilck França and R. Garcia-Patron, *Nature Physics* **17**, 1221 (2021).

[20] D. Wecker, M. B. Hastings, and M. Troyer, *Physical Review A* **92**, 042303 (2015).

[21] J. Stokes, J. Izaac, N. Killoran, and G. Carleo, *Quantum* **4**, 269 (2020).

[22] B. Koczor and S. C. Benjamin, *arXiv preprint arXiv:1912.08660* (2019).

[23] K. M. Nakanishi, K. Fujii, and S. Todo, *Physical Review Research* **2**, 043158 (2020).

[24] J. M. Kübler, A. Arrasmith, L. Cincio, and P. J. Coles, *Quantum* **4**, 263 (2020).

[25] A. Gu, A. Lowe, P. A. Dub, P. J. Coles, and A. Arrasmith, *arXiv preprint arXiv:2108.10434* (2021).

[26] R. Sweke, F. Wilde, J. J. Meyer, M. Schuld, P. K. Fährmann, B. Meynard-Piganeau, and J. Eisert, *Quantum* **4**, 314 (2020).

[27] A. Arrasmith, L. Cincio, R. D. Somma, and P. J. Coles, *arXiv preprint arXiv:2004.06252* (2020).

[28] K. Beer, D. Bondarenko, T. Farrelly, T. J. Osborne, R. Salzmann, D. Scheiermann, and R. Wolf, *Nature Communications* **11**, 808 (2020).

[29] J. Romero, J. P. Olson, and A. Aspuru-Guzik, *Quantum Science and Technology* **2**, 045001 (2017).

[30] R. LaRose, A. Tikku, É. O’Neel-Judy, L. Cincio, and P. J. Coles, *npj Quantum Information* **5**, 1 (2019).

[31] M. Cerezo, K. Sharma, A. Arrasmith, and P. J. Coles, *npj Quantum Information* **8**, 1 (2022).

[32] M. Schuld, A. Bocharov, K. M. Svore, and N. Wiebe, *Physical Review A* **101**, 032308 (2020).

[33] K. Mitarai, M. Negoro, M. Kitagawa, and K. Fujii, *Physical Review A* **98**, 032309 (2018).

[34] M. Schuld, V. Bergholm, C. Gogolin, J. Izaac, and N. Killoran, *Physical Review A* **99**, 032331 (2019).

[35] L. Balles, J. Romero, and P. Hennig, in *Proceedings of the Thirty-Third Conference on Uncertainty in Artificial Intelligence (UAI)* (2017) pp. 410–419.

[36] P. D. Johnson, J. Romero, J. Olson, Y. Cao, and A. Aspuru-Guzik, *arXiv preprint arXiv:1711.02249* (2017).

[37] J. Gibbs, Z. Holmes, M. C. Caro, N. Ezzell, H.-Y. Huang, L. Cincio, A. T. Sornborger, and P. J. Coles, *arXiv preprint arXiv:2204.10269* (2022).

[38] C. Cirstoiu, Z. Holmes, J. Iosue, L. Cincio, P. J. Coles, and A. Sornborger, *npj Quantum Information* **6**, 1 (2020).

[39] M. C. Caro, H.-Y. Huang, N. Ezzell, J. Gibbs, A. T. Sornborger, L. Cincio, P. J. Coles, and Z. Holmes, *arXiv preprint arXiv:2204.10268* (2022).

[40] M. H. Gordon, M. Cerezo, L. Cincio, and P. J. Coles, *PRX Quantum* **3**, 030334 (2022).

[41] S. Thanasilp, S. Wang, N. A. Nghiem, P. J. Coles, and M. Cerezo, *arXiv preprint arXiv:2110.14753* (2021).

[42] B. van Opheusden, L. Acerbi, and W. J. Ma, *PLoS computational biology* **16**, e1008483 (2020).

[43] W. Hoeffding, *The Annals of Mathematical Statistics* **19**, 293 (1948).

[44] S. Lloyd, M. Mohseni, and P. Rebentrost, *Nature Physics* **10**, 631 (2014), number: 9 Publisher: Nature Publishing Group.

### Appendix A: Constructing unbiased estimators for the Log-Likelihood

Here we review a conceptually different approach for a shot-frugal optimization of a problem utilizing a Log-likelihood loss function. In particular, we will follow the results in [42]. First, we assume that one is given a dataset of the form  $\mathcal{S} = \{\rho_i, y_i\}_{i=1}^N$  where  $\rho_i$  are quantum states in some domain  $\mathcal{R}$ , and  $y_i$  are discrete real-valued labels from some domain  $\mathcal{Y}$  (e.g., in binary classification  $\mathcal{Y} = \{0, 1\}$ ). Then, we assume that one has a trainable parametrized model  $h_{\theta} : \mathcal{R} \rightarrow \mathcal{Y}$  whose goal is to predict the correct labels. Note that here one cannot directly apply the techniques previously used, as discrete outputs preclude gradient calculations and the use of the results in [24, 35]. Hence, a different shot-frugal method must be employed.

Given  $\mathcal{S}$  and  $h_{\theta}$ , we can define the likelihood function

$$\Pr(\mathcal{S}|\theta) := \prod_{i=1}^N \Pr(y_i|\rho_i; \theta) = \prod_{i=1}^N p_i, \quad (\text{A1})$$

where we have defined  $p_i = \Pr(y_i|\rho_i; \theta)$  and where assumed that the model’s predictions only depend on the current input. From here, we can define the negative log-likelihood loss function as

$$\mathcal{L}(\theta) = -\log(\Pr(\mathcal{S}|\theta)) = -\sum_{i=1}^N \log(p_i). \quad (\text{A2})$$

Interestingly, we can reduce the task of estimating the negative log-likelihood loss function to a Bernoulli sampling problem. Namely, we can assume without loss of generality that the model specifies a stochastic function  $g$  that takes as input  $\rho_i$ , the set of parameters  $\theta$  and outputs a prediction  $y_i$ . That is,  $y_i \sim g(\rho, \theta)$ . Then, we remark that to estimate the loss function, we do not need to know the explicit form of  $g$ , one only needs to know the probability for a random sample  $y$  from the model to match the true label  $y_i$ . Thus, we can convert each

output from the model into a variable

$$x = \begin{cases} 1 & \text{if } y = y_i \\ 0 & \text{otherwise} \end{cases}. \quad (\text{A3})$$

By construction,  $x$  follows a Bernoulli distribution with probability  $p_i$ . Thus, we can estimate the log-likelihood loss function by drawing samples  $(x_1, x_2, \dots)$  from said Bernoulli distribution. In most common cases these samples can be thought of as single shot-estimates of the model's output, and our goal is to minimize this number of samples/shots.

In [42] the authors study the task of constructing an unbiased estimator for the negative log-likelihood loss function of Eq. (A2), where they consider a fixed number of shots strategy (one uses  $s$  fixed shots) and the inverse binomial sampling (IBS) strategy (where one keeps drawing samples until  $x_k = 1$ ). One can show that the fixed number of shots strategy always leads to a biased estimator of  $\log(p)$  (for the special case when one estimates  $p = \sum_{k=1}^s x_k/s$ , the estimator can even be infinitely biased if  $x_k$  is always zero!). On the other hand, using IBS one can define the estimator

$$\widehat{\mathcal{L}}(\boldsymbol{\theta}) = \begin{cases} 0 & \text{for } K=1 \\ \sum_{k=1}^{K-1} \frac{1}{k} & \text{for } K > 1 \end{cases}, \quad (\text{A4})$$

where  $K$  is the number of shots needed to obtain a result  $x_k = 1$ . Notably, it can be shown that  $E[\widehat{\mathcal{L}}(\boldsymbol{\theta})] = \mathcal{L}(\boldsymbol{\theta})$ , meaning that IBS provides an unbiased estimator for the negative log-likelihood loss function (see [42]). The number of shots that IBS takes trial with probability  $p_i$  is geometrically distributed with mean  $1/p_i$ . For instance, if  $p_i \geq 0.5$ , IBS can require only 1 or 2 shots, while it will assign more shots to cases of small  $p_i$ . We also remark that in Ref. [42] the authors derive the estimator variance, and show that is minimal and non-diverging if  $p_i \rightarrow 0$ . This method then provides a shot-frugal, minimum-variance, unbiased estimator for the negative log-likelihood loss function which can be combined with gradient-free optimizers to train the parameters in the mode. We further refer the reader to Ref. [42] for a discussion on how to extend these results to models  $\ell_{\boldsymbol{\theta}}$  with continuous outputs.

## Appendix B: Using U-statistics to construct unbiased estimators

The theory of U-statistics was initially introduced by Hoeffding in the late 1940s [43] and has a wide range of applications. U-statistics presents a methodology on how to use observations of estimable parameters to construct minimum variance unbiased estimates of more complex functions of the estimable parameters. In the cases we are interested in here, the estimable parameter is the expectation value of an observable used when constructing the cost function.

First, let  $\mathcal{P}$  be a family of probability measures on an arbitrary measurable space. A probability measure is defined as the probability of some particular event occurring. Let  $\Theta(P)$  be a real-valued function defined for  $P \in \mathcal{P}$ . We define  $\Theta(P)$  to be an estimable parameter within  $\mathcal{P}$  if there exists some integer  $r$  for which an unbiased estimator of  $\Theta(P)$  exists constructed from  $r$  independent identically distributed (i.i.d) random variables according to  $P$ . In other words, there is a measurable, real-valued function  $h(x_1, \dots, x_r)$  such that

$$\mathbb{E}_P[h(x_1, \dots, x_r)] = \Theta(P) \quad \forall P \in \mathcal{P}, \quad (\text{B1})$$

where  $x_1, \dots, x_m$  are the i.i.d with distribution  $P$ . The degree of  $\Theta(P)$  is defined to be the smallest integer  $r$  with this property.

The function  $h(x_1, \dots, x_r)$  is also referred to as the kernel. This function can be symmetrized by summing over  $h$  with all possible permutations of the random variables  $x_1, \dots, x_r$ . That is, the symmetrized version of the kernel can be defined as follows

$$h^*(x_1, \dots, x_r) = \frac{1}{r!} \sum_{\sigma \in S_r} h(x_{\sigma(1)}, \dots, x_{\sigma(r)}), \quad (\text{B2})$$

where the summation is over  $S_r$ , the group of all permutations of a set of  $r$  indexes. With these definitions in hand, we can construct a U-statistic.

Assume we are given  $s > r$  samples from  $\mathcal{P}$ , the U-statistic with kernel  $h$  is defined as,

$$U_s = U(h) = \frac{1}{\binom{s}{r}} \sum_{C_r^s} h^*(x_{i_1}, \dots, x_{i_r}), \quad (\text{B3})$$

where the summation is over the set  $C_r^s$  of  $\binom{s}{r} = \frac{s!}{r!(s-r)!}$  combinations of  $r$  integers chosen from  $(1, 2, \dots, s)$ .

As the U-statistic defined above is simply an average of unbiased estimates for  $\Theta(P)$  and is therefore also an unbiased estimator for  $\Theta(P)$ . Indeed, it can be proved that the U-statistic is the best-unbiased estimate for  $\Theta(P)$  in that it has the smallest variance.

### 1. Examples of U-statistics

Now that we have introduced U-statistics we can build intuition with some simple examples. Firstly, let us consider how to construct an unbiased estimate for the mean  $\mu$ . Recall we first need to construct a kernel function  $h$ . In this case, the kernel will have degree 1 and can simply be defined as  $h(x_i) = x_i$ . Note this has the required property that  $\mathbb{E}[h(x_i)] = \mathbb{E}[x_i] = \mu$ . Furthermore, in this case  $h^*(x_i) = h(x_i)$ , i.e. the kernel is already symmetric. Therefore, the U-statistic for  $\mu$  constructed from  $s$

samples can be written as,

$$U_s(h) = \frac{1}{\binom{s}{1}} \sum_{C_1^s} x_i \quad (\text{B4})$$

$$= \frac{1}{s} \sum_{i=1}^s x_i, \quad (\text{B5})$$

which is the well-known formula of the sample average, and which is used to estimate the mean.

Similarly, we can consider how to construct a U-statistic for the variance. Consider the kernel  $h(x_i, x_j) = x_i^2 - x_i x_j$ . This kernel has degree 2. Note once again, this has the required property that  $\mathbb{E}[h(x_i, x_j)] = \mathbb{E}(x_i^2) - \mathbb{E}(x_i x_j) = \mathbb{E}(x_i^2) - \mu^2$ , which is the definition of the variance. Then this needs to be symmetrized,  $h^* = 1/2(x_i^2 - x_i x_j + x_j^2 - x_j x_i) = 1/2(x_i^2 + x_j^2 - 2x_i x_j) = 1/2(x_i - x_j)^2$ . Therefore, the U-statistic for the variance constructed with the kernel  $h$  using  $s$  samples is,

$$U_s(h) = \frac{1}{\binom{s}{2}} \sum_{C_2^s} \frac{(x_i - x_j)^2}{2} \quad (\text{B6})$$

$$= \frac{1}{s(s-1)} \sum_{C_2^s} (x_i - x_j)^2, \quad (\text{B7})$$

which corresponds to the sample variance.

As a final example let's consider constructing an estimator of the function  $\Theta(P) = P^k$ . Lets define the kernel as  $h(x_1, \dots, x_k) = \prod_i x_i$ , which clearly fulfills the properties we require. This kernel is already symmetric such that  $h(x_1, \dots, x_k) = h^*(x_1, \dots, x_m)$ . Therefore, the U-statistic can be defined as,

$$U_s(h) = \frac{1}{\binom{s}{k}} \sum_{C_k^s} \prod_i x_i. \quad (\text{B8})$$

This is exactly the form of estimators that we require in the main text. For our purposes, the number of samples  $s$  is a random variable that is constructed probabilistically.

The framework of U-statistics is perfectly suited to our purposes. Indeed, it provides a recipe to construct unbiased estimators for polynomial functions of variables with some probability distribution. Furthermore, these estimators have the smallest possible variance. Finally, the framework shows the minimum number of shots necessary to estimate a given term through the degree of the kernel.

### Appendix C: Variance of estimator for linear loss functions

In section IV A 2, we constructed an unbiased estimator for the loss function. Denoting the latter  $\hat{\mathcal{L}}$ , we proved in Prop. 1 that  $\mathbb{E}[\hat{\mathcal{L}}] = \mathcal{L}(\theta)$ . We also obtained the variance of such an estimator, which can be used to compare different shot-allocation strategies, similar to [27] for the

Rosalin optimizer designed for VQE. However, we use elements in the variance calculation in Appendix E to construct an unbiased estimator for the MSE loss function.

**Proposition 5.** *The variance of  $\hat{\mathcal{L}}$  is*

$$\begin{aligned} \text{Var}(\hat{\mathcal{L}}) &= \sum_{i,j} \frac{q_{i,j}^2}{\mathbb{E}[s_{i,j}]} \sigma_{i,j}^2 \\ &+ \sum_{i,j,i',j'} \frac{q_{i,j} q_{i',j'}}{\mathbb{E}[s_{i,j}] \mathbb{E}[s_{i',j'}]} \\ &\quad [\langle h_{i,j}(\theta) \rangle \langle h_{i',j'}(\theta) \rangle \text{Cov}[s_{i,j}, s_{i',j'}]], \end{aligned} \quad (\text{C1})$$

where  $\sigma_{i,j}^2 = \langle h_{i,j}^2(\theta) \rangle - (\langle h_{i,j}(\theta) \rangle)^2$  and The covariance is defined as

$$\text{Cov}[s_{i,j}, s_{i',j'}] = \mathbb{E}[s_{i,j} s_{i',j'}] - \mathbb{E}[s_{i,j}] \mathbb{E}[s_{i',j'}]. \quad (\text{C2})$$

*Proof.* Using the definition of the variance and recalling that  $\mathbb{E}[\hat{\mathcal{L}}] = \mathcal{L}(\theta)$ , we can write

$$\begin{aligned} \text{Var}(\hat{\mathcal{L}}) &= \mathbb{E}[\hat{\mathcal{L}}^2] - \mathbb{E}[\hat{\mathcal{L}}]^2 \\ &= \sum_{i,j,i',j'} \frac{q_{i,j} q_{i',j'}}{\mathbb{E}[s_{i,j}] \mathbb{E}[s_{i',j'}]} \mathbb{E}[\sum_{k=1}^{s_{i,j}} \sum_{k'=1}^{s_{i',j'}} r_{i,j,k} r_{i',j',k'}] \\ &\quad - \mathcal{L}(\theta)^2. \end{aligned} \quad (\text{C3})$$

We can then decompose the first term into its constituent parts where indices in the sum are different and equal

$$\begin{aligned} &\sum_{i,j,i',j'} \frac{q_{i,j} q_{i',j'}}{\mathbb{E}[s_{i,j}] \mathbb{E}[s_{i',j'}]} \mathbb{E}[\sum_{k=1}^{s_{i,j}} \sum_{k'=1}^{s_{i',j'}} r_{i,j,k} r_{i',j',k'}] \\ &= \sum_{i,j,i',j',i \neq i', j \neq j'} \frac{q_{i,j} q_{i',j'}}{\mathbb{E}[s_{i,j}] \mathbb{E}[s_{i',j'}]} E_1 \\ &+ \sum_{i,j,i',j',i=i', j \neq j'} \frac{q_{i,j} q_{i',j'}}{\mathbb{E}[s_{i,j}] \mathbb{E}[s_{i',j'}]} E_2 \\ &+ \sum_{i,j,i',j',i \neq i', j=j'} \frac{q_{i,j} q_{i',j}}{\mathbb{E}[s_{i,j}] \mathbb{E}[s_{i',j}]} E_3 \\ &+ \sum_{i,j} \frac{q_{i,j}^2}{\mathbb{E}^2[s_{i,j}]} E_4, \end{aligned} \quad (\text{C4})$$

where

$$E_1 = \mathbb{E}[\sum_{k=1}^{s_{i,j}} \sum_{k'=1}^{s_{i',j'}} r_{i,j,k} r_{i',j',k'}] \quad (\text{C5})$$

$$E_2 = \mathbb{E}[\sum_{k=1}^{s_{i,j}} \sum_{k'=1}^{s_{i',j'}} r_{i,j,k} r_{i',j',k'}] \quad (\text{C6})$$

$$E_3 = \mathbb{E}[\sum_{k=1}^{s_{i,j}} \sum_{k'=1}^{s_{i',j}} r_{i,j,k} r_{i',j,k'}] \quad (\text{C7})$$

$$E_4 = \mathbb{E}[\sum_{k=1}^{s_{i,j}} \sum_{k'=1}^{s_{i,j}} r_{i,j,k} r_{i,j,k'}]. \quad (\text{C8})$$

Note that since  $\frac{1}{\mathbb{E}[s_{i,j}]} \sum_{k=1}^{s_{i,j}} r_{i,j,k}$  is an unbiased estimator of  $\langle h_{i,j}(\boldsymbol{\theta}) \rangle$ , by application of Wald's equation we obtain □

$$E_1 = \langle h_{i,j}(\boldsymbol{\theta}) \rangle \langle h_{i',j'}(\boldsymbol{\theta}) \rangle \mathbb{E}[s_{i,j} s_{i',j'}] \quad (\text{C9})$$

$$E_2 = \langle h_{i,j}(\boldsymbol{\theta}) \rangle \langle h_{i',j'}(\boldsymbol{\theta}) \rangle \mathbb{E}[s_{i,j} s_{i',j'}] \quad (\text{C10})$$

$$E_3 = \langle h_{i,j}(\boldsymbol{\theta}) \rangle \langle h_{i',j'}(\boldsymbol{\theta}) \rangle \mathbb{E}[s_{i,j} s_{i',j'}]. \quad (\text{C11})$$

For the case  $\{i = i', j = j'\}$ , care needs to be taken when expanding the sum as the terms are not independent when  $k' = k$ . Therefore, we decompose the sum into the cases where  $k \neq k'$  and  $k = k'$ ,

$$\begin{aligned} E_4 &= \mathbb{E}\left[\sum_{k=1}^{s_{i,j}} \sum_{k'=1, k' \neq k}^{s_{i,j}} r_{i,j,k} r_{i,j,k'}\right] + \mathbb{E}\left[\sum_{k=1}^{s_{i,j}} r_{i,j,k}^2\right] \\ &= E_{4,1} + E_{4,2}. \end{aligned} \quad (\text{C12})$$

Once again, using Wald's equation, but taking into account that the number of entries for the  $k \neq k'$  is  $(s_{i,j} - 1)$ , one can show

$$E_{4,1} = \mathbb{E}[s_{i,j}(s_{i,j} - 1)] (\langle h_{i,j}(\boldsymbol{\theta}) \rangle)^2. \quad (\text{C13})$$

Finally, using the fact that  $\frac{1}{\mathbb{E}[s_{i,j}]} \sum_{k=1}^{s_{i,j}} r_{i,j,k}^2$  is an unbiased estimator of  $\langle h_{i,j}^2(\boldsymbol{\theta}) \rangle$ , we obtain:

$$E_{4,2} = \mathbb{E}[s_{i,j}] \langle h_{i,j}^2(\boldsymbol{\theta}) \rangle. \quad (\text{C14})$$

Bringing all the terms together and writing the expressions containing  $E_1$ ,  $E_2$  and  $E_3$  in a single sum, we obtain

$$\begin{aligned} \text{Var}(\widehat{\mathcal{L}}) &= \sum_{i,j} \frac{q_{i,j}^2}{\mathbb{E}^2[s_{i,j}]} \mathbb{E}[s_{i,j}(s_{i,j} - 1)] (\langle h_{i,j}(\boldsymbol{\theta}) \rangle)^2 \\ &+ \sum_{i,j} \frac{q_{i,j}^2}{\mathbb{E}^2[s_{i,j}]} \mathbb{E}[s_{i,j}] \langle h_{i,j}^2(\boldsymbol{\theta}) \rangle - \mathcal{L}(\boldsymbol{\theta})^2 \\ &+ \sum_{\substack{i,j,i',j' \\ i \neq i' \text{ if } j=j' \\ j \neq j' \text{ if } i=i'}} \frac{q_{i,j} q_{i',j'}}{\mathbb{E}[s_{i,j}] \mathbb{E}[s_{i',j'}]} \langle h_{i,j}(\boldsymbol{\theta}) \rangle \langle h_{i',j'}(\boldsymbol{\theta}) \rangle \\ &\mathbb{E}[s_{i,j} s_{i',j'}]. \end{aligned} \quad (\text{C15})$$

Expanding the second term and factorizing leads to,

$$\begin{aligned} \text{Var}(\widehat{\mathcal{L}}) &= \sum_{i,j} \frac{q_{i,j}^2}{\mathbb{E}[s_{i,j}]} [\langle h_{i,j}^2(\boldsymbol{\theta}) \rangle - \langle h_{i,j}(\boldsymbol{\theta}) \rangle^2] \\ &+ \sum_{i,j,i',j'} \frac{q_{i,j} q_{i',j'}}{\mathbb{E}[s_{i,j}] \mathbb{E}[s_{i',j'}]} \mathbb{E}[s_{i,j} s_{i',j'}] \\ &\langle h_{i,j}(\boldsymbol{\theta}) \rangle \langle h_{i',j'}(\boldsymbol{\theta}) \rangle - \mathcal{L}(\boldsymbol{\theta})^2. \end{aligned} \quad (\text{C16})$$

Which we can then rewrite to prove the desired result,

$$\begin{aligned} \text{Var}(\widehat{\mathcal{L}}) &= \sum_{i,j} \frac{q_{i,j}^2}{\mathbb{E}[s_{i,j}]} \sigma_{i,j}^2 \\ &+ \sum_{i,j,i',j'} \frac{q_{i,j} q_{i',j'}}{\mathbb{E}[s_{i,j}] \mathbb{E}[s_{i',j'}]} \langle h_{i,j}(\boldsymbol{\theta}) \rangle \langle h_{i',j'}(\boldsymbol{\theta}) \rangle \text{Cov}[s_{i,j}, s_{i',j'}]. \end{aligned} \quad (\text{C17})$$

## Appendix D: Constructing an unbiased estimators of loss functions with polynomial dependence

We consider here how to construct estimators for the loss functions with an arbitrary polynomial of degree  $D$  in the observables, expressed as:

$$\mathcal{L}(\boldsymbol{\theta}) = \sum_{\mathbf{i}} \sum_{z=0}^D p_{\mathbf{i},z} a_z \left( \sum_j c_{i,j} \langle h_{i,j}(\boldsymbol{\theta}) \rangle \right)^z \quad (\text{D1})$$

$$= \sum_{\mathbf{i},z} p_{\mathbf{i},z} \left( \sum_j c_{i,j} \langle h_{i,j}(\boldsymbol{\theta}) \rangle \right)^z, \quad (\text{D2})$$

where  $p_{\mathbf{i},z} = p_{\mathbf{i}} a_z$ ,  $a_z \in \mathbb{R}$ . Using the multinomial theorem, given  $J$  hamiltonian terms, we can develop the power of the sum into a sum of  $\binom{z+J-1}{J}$  terms:

$$\mathcal{L}(\boldsymbol{\theta}) = \sum_{\mathbf{i},z} p_{\mathbf{i},z} \sum_{\substack{b_1+b_2+\dots+b_J=z}} \binom{z}{b_1, b_2, \dots, b_J} \prod_j (c_{i,j} \langle h_{i,j}(\boldsymbol{\theta}) \rangle)^{b_j}, \quad (\text{D3})$$

where  $\binom{z}{b_1, b_2, \dots, b_J} = \frac{z!}{b_1! b_2! \dots b_J!}$  and  $b_j$  are non-negative integers. First, one can generalize the formula for an unbiased estimator of  $(\langle h_{i,j}(\boldsymbol{\theta}) \rangle)^{b_j}$ :

$$\widehat{\mathcal{E}}_{i,j} = \frac{1}{\mathbb{E}[\prod_{t=0}^{b_j-1} (s_{i,j} - t)]} \prod_{\alpha=1}^{b_j} \left( \sum_{k_\alpha}^{s_{i,j}} r_{i,j,k_\alpha} \right). \quad (\text{D4})$$

The case  $b_j = 2$  was elaborated in Eq.E11.

For  $\prod_j \langle h_{i,j}(\boldsymbol{\theta}) \rangle^{b_j}$ , we define an unbiased estimator as:

$$\widehat{\mathcal{R}}_{i,z} = \frac{1}{\mathbb{E}[\prod_j \prod_{t=0}^{b_j-1} (s_{i,j} - t)]} \prod_j \prod_{\alpha=1}^{b_j} \left( \sum_{k_{j,\alpha}}^{s_{i,j}} r_{i,j,k_{j,\alpha}} \right). \quad (\text{D5})$$

This involves estimating  $\binom{z+J-1}{J} \prod_{j=1}^J b_j$  terms to construct the following estimator for the cost:

$$\widehat{\mathcal{L}} = \sum_{\mathbf{i},z} p_{\mathbf{i},z} \sum_{\substack{b_1+b_2+\dots+b_J=z}} \binom{z}{b_1, b_2, \dots, b_J} \left[ \prod_j (c_{i,j})^{b_j} \right] \widehat{\mathcal{R}}_{i,z}, \quad (\text{D6})$$

and establish the following proposition.

**Proposition 6.**  $\widehat{\mathcal{L}}$  is an unbiased estimator for  $\mathcal{L}$ .

*Proof.* Taking the expectation over the above definitions:

$$\begin{aligned} \mathbb{E}[\widehat{\mathcal{L}}] &= \sum_{\mathbf{i},z} p_{\mathbf{i},z} \sum_{\substack{b_1+b_2+\dots+b_J=z}} \binom{z}{b_1, b_2, \dots, b_J} \left[ \prod_j (c_{i,j})^{b_j} \right] \mathbb{E}[\widehat{\mathcal{R}}_{i,z}] \\ &= \sum_{\mathbf{i},z} p_{\mathbf{i},z} \sum_{b_1+b_2+\dots+b_J=z} \binom{z}{b_1, b_2, \dots, b_J} \left[ \prod_j (c_{i,j})^{b_j} \right] \\ &\frac{1}{\mathbb{E}[\prod_j \prod_{t=0}^{b_j-1} (s_{i,j} - t)]} \mathbb{E}[\prod_j \prod_{\alpha=1}^{b_j} \sum_{k_{j,\alpha}}^{s_{i,j}} r_{i,j,k_{j,\alpha}}]. \end{aligned} \quad (\text{D7})$$

Using again Wald's inequality, we obtain:

$$\mathbb{E}\left[\prod_j \prod_{\alpha=1}^{b_j} \sum_{k_j, \alpha}^{s_{i,j}} r_{i,j,k_j, \alpha}\right] = \mathbb{E}\left[\prod_j \prod_{t=0}^{b_j-1} (s_{i,j} - t)\right] \prod_j \langle h_{i,j}(\boldsymbol{\theta}) \rangle^{b_j}.$$

for which, when replacing in the previous equation, we obtain  $\mathbb{E}[\widehat{\mathcal{L}}] = \mathcal{L}(\boldsymbol{\theta})$ .  $\square$

## Appendix E: Constructing an unbiased estimator for the Mean Squared Error Loss Function

The mean squared error (MSE) is widely used in machine learning. Here, we show how to construct an unbiased estimator for the MSE as an illustrative example of previously-studied polynomial function of the observables in Appendix D4. We start by considering a simplified case for the MSE, where we do not decompose each measurement operator  $H_i$  into a linear combination of Hermitian matrices. We then tackle the more complex case where  $H_i$  is of the form presented in Eq. (11).

### 1. Step 1: Simplified MSE

Let  $\{\rho_i, y_i\}$  be a dataset of  $N$  points, where the task consists of minimizing the following cost:

$$\mathcal{L}_{MSE} = \sum_i p_i (y_i - \langle h_i(\boldsymbol{\theta}) \rangle)^2, \quad (\text{E1})$$

This corresponds to a weighted MSE cost. The most common case is when  $p_i = 1/N$ . We will take a step-by-step approach to construct unbiased estimators for such cost.

If we expand the brackets, we obtain,

$$\mathcal{L}_{MSE} = \sum_i p_i (y_i^2 - 2y_i \langle h_i(\boldsymbol{\theta}) \rangle + \langle h_i(\boldsymbol{\theta}) \rangle^2).$$

We denote  $\widehat{\mathcal{E}}_i$  as the estimator for  $\langle h_i(\boldsymbol{\theta}) \rangle$ ,  $\widehat{\mathcal{Q}}_i$  as the estimator for  $\langle h_i(\boldsymbol{\theta}) \rangle^2$ .

$$\widehat{\mathcal{E}}_i = \frac{1}{\mathbb{E}[s_i]} \sum_{k=1}^{s_i} r_{i,k}, \quad (\text{E2})$$

$$\widehat{\mathcal{Q}}_i = \frac{1}{\mathbb{E}[s_i(s_i - 1)]} \sum_{k=1}^{s_i} \sum_{k'=1, k' \neq k}^{s_i} r_{i,k} r_{i,k'}, \quad (\text{E3})$$

where  $\widehat{\mathcal{E}}_i$  is an unbiased estimator which is straightforward from Wald's equation.  $\widehat{\mathcal{Q}}_i$  is also an unbiased estimator for  $\langle h_i(\boldsymbol{\theta}) \rangle^2$ , following the reasoning used in Eq. (C13) in the proof of Prop. 5. Then, defining  $\widehat{\mathcal{L}}_{MSE}$  as

$$\widehat{\mathcal{L}}_{MSE} = \sum_i p_i (y_i^2 - 2y_i \widehat{\mathcal{E}}_i + \widehat{\mathcal{Q}}_i), \quad (\text{E4})$$

we establish the following proposition.

**Proposition 7.**  $\widehat{\mathcal{L}}_{MSE}$  is an unbiased estimator for  $\mathcal{L}_{MSE}$ .

*Proof.* Using the definitions in Eq. (E2) and Eq. (E4) we can write

$$\begin{aligned} \mathbb{E}[\widehat{\mathcal{L}}_{MSE}] &= \mathbb{E}\left[\sum_i p_i y_i^2\right] - 2\mathbb{E}\left[\sum_i y_i p_i \frac{1}{\mathbb{E}[s_i]} \sum_{k=1}^{s_i} r_{i,k}\right] \\ &+ \mathbb{E}\left[\sum_i p_i \frac{1}{\mathbb{E}[s_i(s_i - 1)]} \sum_{k=1}^{s_i} \sum_{k'=1, k' \neq k}^{s_i} r_{i,k} r_{i,k'}\right]. \quad (\text{E5}) \end{aligned}$$

Using Wald's equation we can evaluate this expression to give

$$\begin{aligned} \mathbb{E}[\widehat{\mathcal{L}}_{MSE}] &= \sum_i p_i y_i^2 - 2 \sum_i y_i p_i \frac{\mathbb{E}[s_i]}{\mathbb{E}[s_i]} \langle h_i(\boldsymbol{\theta}) \rangle \\ &+ \sum_i p_i \frac{\mathbb{E}[s_i(s_i - 1)]}{\mathbb{E}[s_i(s_i - 1)]} (\langle h_i(\boldsymbol{\theta}) \rangle)^2. \quad (\text{E6}) \end{aligned}$$

Finally, we can evaluate these terms and factorize, which leads to,

$$\mathbb{E}[\widehat{\mathcal{L}}_{MSE}] = \sum_i p_i (y_i^2 - 2y_i \langle h_i(\boldsymbol{\theta}) \rangle + \langle h_i(\boldsymbol{\theta}) \rangle^2) \quad (\text{E7})$$

$$= \mathcal{L}_{MSE}. \quad (\text{E8})$$

$\square$

This means that we can estimate this simplified MSE loss function by evaluating the first term exactly, sampling from the second term according to the probability distribution  $p_i y_i$  and sampling the third term according to the distribution  $p_i$ .

### 2. Step 2: Full MSE

Now, we assume the Hamiltonians  $H_i$  can be decomposed into many terms such that  $H_i = \sum_j c_{i,j} H_{i,j}$ . Hence the new full MSE cost can be expressed as

$$\begin{aligned} \mathcal{L}'_{MSE} &= \sum_i p_i (y_i - \sum_j c_{i,j} \langle h_{i,j}(\boldsymbol{\theta}) \rangle)^2 \\ &= \sum_i p_i y_i^2 - 2 \sum_{i,j} p_i y_i c_{i,j} \langle h_{i,j}(\boldsymbol{\theta}) \rangle \\ &+ \sum_{i,j,j',j' \neq j} p_i c_{i,j} c_{i,j'} \langle h_{i,j}(\boldsymbol{\theta}) \rangle \langle h_{i,j'}(\boldsymbol{\theta}) \rangle + \sum_{i,j} p_i c_{i,j}^2 \langle h_{i,j}(\boldsymbol{\theta}) \rangle^2. \quad (\text{E9}) \end{aligned}$$

Defining  $\widehat{\mathcal{E}}_{i,j}$  as the estimator for  $T_{i,j}$  as previously,

$$\widehat{\mathcal{E}}_{i,j} = \frac{1}{\mathbb{E}[s_{i,j}]} \sum_{k=1}^{s_{i,j}} r_{i,j,k}.$$

When  $j \neq j'$  we define,

$$\widehat{\mathcal{Q}}_{i,j,j'} = \frac{1}{\mathbb{E}[s_{i,j} s_{i,j'}]} \sum_{k=1}^{s_{i,j}} \sum_{k'=1}^{s_{i,j'}} r_{i,j,k} r_{i,j',k'}, \quad (\text{E10})$$

and finally,

$$\hat{Q}_{i,j} = \frac{1}{\mathbb{E}[s_{i,j}(s_{i,j} - 1)]} \sum_{k=1}^{s_{i,j}} \sum_{k'=1, k' \neq k}^{s_{i,j}} r_{i,j,k} r_{i,j,k'}. \quad (\text{E11})$$

Collecting these terms we define an estimator of the full MSE loss function as

$$\begin{aligned} \hat{\mathcal{L}}'_{MSE} &= \sum_{\mathbf{i}} p_{\mathbf{i}} y_{\mathbf{i}}^2 - 2 \sum_{\mathbf{i}, j} p_{\mathbf{i}} y_{\mathbf{i}} c_{\mathbf{i}, j} \hat{\mathcal{E}}_{\mathbf{i}, j} \\ &+ \sum_{\mathbf{i}, j, j', j' \neq j} p_{\mathbf{i}} c_{\mathbf{i}, j} c_{\mathbf{i}, j'} \hat{Q}_{\mathbf{i}, j, j'} + \sum_{\mathbf{i}, j} p_{\mathbf{i}} c_{\mathbf{i}, j}^2 \hat{Q}_{\mathbf{i}, j}, \end{aligned} \quad (\text{E12})$$

which can be used to establish the following proposition.

**Proposition 8.**  $\hat{\mathcal{L}}'_{MSE}$  is an unbiased estimator for  $\mathcal{L}'_{MSE}$ .

*Proof.* Taking the above definition we can write the fol-

lowing result:

$$\begin{aligned} \mathbb{E}[\hat{\mathcal{L}}'_{MSE}] &= \sum_{\mathbf{i}} p_{\mathbf{i}} y_{\mathbf{i}}^2 - 2 \sum_{\mathbf{i}, j} p_{\mathbf{i}} y_{\mathbf{i}} c_{\mathbf{i}, j} \frac{1}{\mathbb{E}[s_{i,j}]} \mathbb{E}\left[\sum_{k=1}^{s_{i,j}} r_{i,j,k}\right] \\ &+ \sum_{\mathbf{i}, j, j', j' \neq j} p_{\mathbf{i}} c_{\mathbf{i}, j} c_{\mathbf{i}, j'} \frac{1}{\mathbb{E}[s_{i,j} s_{i,j'}]} \mathbb{E}\left[\sum_{k=1}^{s_{i,j}} \sum_{k'=1}^{s_{i,j'}} r_{i,j,k} r_{i,j',k'}\right] \\ &+ \sum_{\mathbf{i}, j} p_{\mathbf{i}} c_{\mathbf{i}, j}^2 \frac{1}{\mathbb{E}[s_{i,j}(s_{i,j} - 1)]} \mathbb{E}\left[\sum_{k=1}^{s_{i,j}} \sum_{k'=1, k' \neq k}^{s_{i,j}} r_{i,j,k} r_{i,j,k'}\right]. \end{aligned} \quad (\text{E13})$$

Given that

$$\mathbb{E}\left[\sum_{k=1}^{s_{i,j}} r_{i,j,k}\right] = \mathbb{E}[s_{i,j}] \langle h_{i,j}(\boldsymbol{\theta}) \rangle \quad (\text{E14})$$

$$\mathbb{E}\left[\sum_{k=1}^{s_{i,j}} \sum_{k'=1}^{s_{i,j'}} r_{i,j,k} r_{i,j',k'}\right] = \mathbb{E}[s_{i,j} s_{i,j'}] \langle h_{i,j}(\boldsymbol{\theta}) \rangle \langle h_{i,j'}(\boldsymbol{\theta}) \rangle \quad (\text{E15})$$

$$\mathbb{E}\left[\sum_{k=1}^{s_{i,j}} \sum_{k'=1, k' \neq k}^{s_{i,j}} r_{i,j,k} r_{i,j,k'}\right] = \mathbb{E}[s_{i,j}(s_{i,j} - 1)] \langle h_{i,j}(\boldsymbol{\theta}) \rangle^2. \quad (\text{E16})$$

We can combine the above expressions to obtain

$$\begin{aligned} \mathbb{E}[\hat{\mathcal{L}}'_{MSE}] &= \sum_{\mathbf{i}} p_{\mathbf{i}} y_{\mathbf{i}}^2 - 2 \sum_{\mathbf{i}, j} p_{\mathbf{i}} y_{\mathbf{i}} c_{\mathbf{i}, j} \langle h_{i,j}(\boldsymbol{\theta}) \rangle \\ &+ \sum_{\mathbf{i}, j, j', j' \neq j} p_{\mathbf{i}} c_{\mathbf{i}, j} c_{\mathbf{i}, j'} \langle h_{i,j}(\boldsymbol{\theta}) \rangle \langle h_{i,j'}(\boldsymbol{\theta}) \rangle \\ &+ \sum_{\mathbf{i}, j} p_{\mathbf{i}} c_{\mathbf{i}, j}^2 \langle h_{i,j}(\boldsymbol{\theta}) \rangle^2 \end{aligned} \quad (\text{E17})$$

$$= \mathcal{L}'_{MSE}. \quad (\text{E18})$$

□

This shows we can estimate the full MSE cost function by evaluating 4 terms. One can be evaluated exactly and the others can be estimated by sampling from their corresponding probability distributions built with the weights as parameters.