

Noisy Bayesian optimization for variational quantum eigensolvers

Giovanni Iannelli^{a,b,c,d,*} and Karl Jansen^a

^a*Deutsches Elektronen-Synchrotron DESY,
Platanenallee 6, 15738 Zeuthen, Germany*

^b*Department of Physics, University of Cyprus,
Panepistimiou Street 1, 2109 Aglantzia, Nicosia, Cyprus*

^c*Institut für Physik, Humboldt-Universität zu Berlin,
Newtonstraße 15, 12489 Berlin, Germany*

^d*Dipartimento di Fisica, Università degli Studi di Roma "Tor Vergata",
Via della Ricerca Scientifica 1, 00133 Rome, Italy*

E-mail: giovanni.iannelli@desy.de, karl.jansen@desy.de

The variational quantum eigensolver (VQE) is a hybrid quantum-classical algorithm used to find the ground state of a Hamiltonian using variational methods. In the context of this Lattice symposium, the procedure can be used to study lattice gauge theories (LGTs) in the Hamiltonian formulation. Bayesian optimization (BO) based on Gaussian process regression (GPR) is a powerful algorithm for finding the global minimum of a cost function, e.g. the energy, with a very low number of iterations using data affected by statistical noise. This work proposes an implementation of GPR and BO specifically tailored to perform VQE on quantum computers already available today.

*The 38th International Symposium on Lattice Field Theory, LATTICE2021 26th-30th July, 2021
Zoom/Gather@Massachusetts Institute of Technology*

*Speaker

1. Introduction

Quantum algorithms have the potential to be exponentially quicker than classical alternatives in many noteworthy scientific applications. Examples are quantum machine learning [1], quantum chemistry [2], and many others [3]. Unfortunately, many of these applications are not yet implementable on current noisy intermediate-scale quantum (NISQ) computers [4] and need to wait until noise sources can be suppressed down to a threshold that makes quantum computers usable in practice or to even build fault-tolerant quantum computers [5].

However, many interesting problems of LGTs can already be studied with NISQ devices [6]. In particular, if LGTs are studied in their Hamiltonian formulation, quantum algorithms do not generally suffer from the sign problem [7, 8]. An important ready-to-use algorithm is the variational quantum eigensolver (VQE) [9], which is a hybrid quantum-classical algorithm for finding the ground (and excited) state of a given Hamiltonian \mathcal{H} using the variational principle. The *quantum part* of VQE deals with measuring the expectation value of the Hamiltonian, i.e. the energy, in a given multi-qubit state, while the *classical part* consists of searching among a family of multi-qubit states generated by a parametrized quantum circuit to find the state that minimizes the energy.

The algorithm proposed in this proceedings is a classical optimizer that aims to find a good approximation of the ground state reducing as much as possible the number of energy measurements. The approach chosen here is known as Bayesian global optimization. Its first application dates back in the 60s [10], while its modern implementations are based on a more recent work [11]. The backbone of this method is Gaussian process regression (GPR), which is an interpolation method based on Bayesian inference of Gaussian processes. It allows us to create predictive models of black-box functions using a limited amount of (noisy) data. At each optimization iteration, this model is used to determine a set of parameters presumably close to the global minimum point. This step is performed following a procedure called acquisition function optimization.

The algorithm proposed here to optimize the energy differs from the other alternatives commonly used in VQE as it uses not only the estimated values of the energy, but also the values of their statistical errors. The motivation is to lower the number of quantum measurements at each step: the procedure is well defined even for imprecise energy measurements as long as their errors is approximately Gaussian due to the central limit theorem. Results of this algorithm are compared to other commonly chosen alternatives using simulators of noisy devices.

2. Quantum expectation estimation

Given a Hamiltonian \mathcal{H} , it first needs to be written as a polynomial of sigma matrices:

$$H = \sum_{i\gamma} h_{\gamma}^i \sigma_{\gamma}^i + \sum_{ij\gamma\delta} h_{\gamma\delta}^{ij} \sigma_{\gamma}^i \otimes \sigma_{\delta}^j + \dots \quad (1)$$

where the h s are real parameters, Latin indices identify the qubit on which sigma matrices are acting, and Greek indices the sigma matrices coordinates. The quantum expectation estimation (QEE) algorithm [9] computes the expectation value of the energy $E_{\psi} \equiv \langle \psi | \mathcal{H} | \psi \rangle$ for any input multi-qubit state $|\psi\rangle$ with a possible quantum advantage with respect to equivalent classical

approaches. Furthermore, QEE is already implementable in NISQ devices as the computation of E_ψ can be decomposed in many short quantum programs, therefore reducing the impact of quantum noise.

However, due to the probabilistic nature of quantum measurements, we have only access to a stochastic variable that estimates E_ψ (see for example [12] for the Qiskit implementation). In order to get a precise estimation, it is possible to perform multiple independent measurements, also called *shots*, and use their mean as an estimator of E_ψ . Let us consider the case of measurements in the absence of quantum noise. As the number of shots goes to infinity, the central limit theorem tells us that the mean converges to a Gaussian centered in E_ψ whose variance is estimated by the standard error of the mean. Therefore, given a number S of independent shots, calling $E_\psi^1, \dots, E_\psi^S$ the energy measurements of each single shot, it is possible to measure the following energy estimator on a quantum computer:

$$\hat{E}_\psi^S \equiv \frac{1}{S} \sum_{m=1}^S E_\psi^m \quad (2)$$

$$\text{Var}[\hat{E}_\psi^S] = \frac{1}{S(S-1)} \sum_{m=1}^S (E_\psi^m - \hat{E}_\psi^S)^2 \quad (3)$$

It is important to emphasize that there is a difference between the *statistical noise* and the *quantum noise*. The statistical noise is the (approximately) Gaussian deviation of \hat{E}_ψ^S caused by the probabilistic nature of quantum measurements, while quantum noise is the deviation caused by the imperfections of real quantum devices. The impact of quantum noise to QEE is to add a BIAS to the estimator of Eq. (2). This BIAS can be significantly reduced using error mitigation techniques (a comparison can, e.g., be found in Section 5.1 of [13]).

In the rest of these proceedings, for simplifying the notation, the number of shots S will be omitted. Since on the quantum computer a parametrized quantum circuit is employed with parameters θ_α , we will denote the parametrized energy $E(\theta_\alpha)$ with error $\Delta E(\theta_\alpha)$ with $E(\theta_\alpha) \equiv \hat{E}^S(\theta_\alpha)$ and $\Delta E(\theta_\alpha) \equiv \sqrt{\text{Var}[\hat{E}^S(\theta_\alpha)]}$.

3. Variational quantum eigensolver

The objective Hamiltonian \mathcal{H} needs first to be written in the form of, e.g., Eq. (1). Then, a family of qubit states $|\psi(\theta_\alpha)\rangle$ is introduced with a d -dimensional parameter set θ_α . This can be achieved by applying a parametrized quantum circuit $U(\theta_\alpha)$ on a fixed initial multi-qubit state, usually chosen to be $|0 \dots 0\rangle$:

$$|\psi(\theta_\alpha)\rangle \equiv U(\theta_\alpha) |0 \dots 0\rangle \quad (4)$$

This state parametrization allows us to define a parametrized energy:

$$E(\theta_\alpha) \equiv \langle \psi(\theta_\alpha) | \mathcal{H} | \psi(\theta_\alpha) \rangle \quad (5)$$

which can be evaluated using the QEE algorithm for any value of θ_α .

Then, the VQE consists of approximating the ground state $|\psi_{\min}\rangle$ performing the following optimization:

$$\min_{\theta_\alpha} E(\theta_\alpha) \geq \langle \psi_{\min} | \mathcal{H} | \psi_{\min} \rangle = E_{\min} \quad (6)$$

It is important to note that most of the optimizers commonly used within VQE take as input only measurements of the estimator in Eq. (2) and not of its error in Eq. (3). In some cases they rely on performing energy measurements that are precise enough to be considered (almost) exact, which means choosing the number of shots S to be large enough so that the error in Eq. (3) can be neglected. This is done, for example, in the original VQE paper [9] using the Nelder-Mead [14] optimizer, as well as in a well-performing recently published optimizer algorithm [15] specifically tailored for VQE. On the other hand, some algorithms leverage on the statistical noise of input measurements to escape local minima, as for example in the SPSA [16] optimizer. The algorithm proposed in this proceedings differs from the considered alternatives as it uses both the estimator of Eq. (2) and its error of Eq. (3).

4. Bayesian optimization

In the Bayesian optimization approach one first needs an initialization step with a sequence of energy measurements obtained with circuit parameters chosen (quasi) randomly. After this initial step, the core of the algorithm consists of two building blocks. At each iteration, it first uses GPR to create a predictive model of the parametrized energy in Eq. (5) using the information of the previous energy measurements. Then, this predictive model is used to define an acquisition function that assigns to each circuit parameter values a positive score. The parameters that maximize the acquisition function, i.e. the one with the highest score, will then be chosen for the next energy measurement.

4.1 Gaussian process regression

A Gaussian process (GP) f maps a set of n d -dimensional circuit parameters $\theta_{1\alpha}, \dots, \theta_{n\alpha}$ to n stochastic variables f_1, \dots, f_n distributed as:

$$p(f_i) = \det(2\pi K)^{-\frac{1}{2}} \exp\left(-\frac{1}{2} \sum_{ij} (f_i - \mu_i)(K^{-1})_{ij}(f_j - \mu_j)\right)$$

where $\mu_i \equiv \mu(\theta_{i\alpha})$, $K_{ij} \equiv k(\theta_{i\alpha}, \theta_{j\alpha})$, $\mu(\theta_\alpha)$ is the GP mean function and $k(\theta_\alpha, \theta'_\alpha)$ is the GP covariance function.

Given a set of energy measurements E_1, \dots, E_n obtained with the corresponding circuit parameter values $\theta_{1\alpha}, \dots, \theta_{n\alpha}$ with their respective measurement errors $\Delta E_1, \dots, \Delta E_n$, Gaussian process regression¹ (GPR) [17] is a procedure that finds a GP whose mean function $\mu(\theta_\alpha)$ interpolates the unknown energy functions $E(\theta)$. For details about GPR see, e.g., the textbook [18]. For what concerns our presentation, we report the analytical result for the mean and covariance functions of the posterior obtained with GPR:

$$\begin{cases} \mu(\theta_\alpha | E_i) = \mu(\theta_\alpha) + \sum_{ij} k(\theta_\alpha, \theta_{i\alpha})(\tilde{K}^{-1})_{ij}(E_j - \mu(\theta_{j\alpha})) \\ k(\theta_\alpha, \theta'_\alpha | E_i) = k(\theta_\alpha, \theta'_\alpha) - \sum_{ij} k(\theta_\alpha, \theta_{i\alpha})(\tilde{K}^{-1})_{ij}k(\theta_{j\alpha}, \theta'_\alpha) \end{cases} \quad (7)$$

¹Also known as kriging in geostatistics.

where $\tilde{K}_{ij} \equiv k(\theta_{i\alpha}, \theta_{j\alpha}) + \Delta E_i^2 \delta_{ij}$. The posterior mean $\mu(\theta_\alpha | E_i)$ is our surrogate model of the unknown parametrized energy, while $k(\theta_\alpha, \theta'_\alpha | E_i)$ is an estimation of its Gaussian covariance. We emphasize that the measurement errors evaluated with Eq. (3) are used in the evaluation of \tilde{K}_{ij} , and this formula is exact in the case of Gaussian errors, which is asymptotically true as the number of shots grows.

The choice of the prior mean $\mu(\theta)$ and the prior covariance function $k(\theta_\alpha, \theta'_\alpha)$ is subjective, as it usually happens in the context of Bayesian inference. Their choice has an impact on the geometry of the posterior GP of Eq. (7). The possibility of using different $\mu(\theta)$ and $k(\theta_\alpha, \theta'_\alpha)$ can then be an advantage as it can be used to impose certain properties that are motivated from the physics of the considered problem. For example, in the noiseless case, the parametrized energy $E(\theta_\alpha)$ of Eq. (5) is in C^∞ when using commonly chosen quantum circuits [15], and selecting $\mu(\theta_\alpha), k(\theta_\alpha, \theta'_\alpha) \in C^\infty$ will impose this property on the posterior mean of Eq. (7). A common choice is to set $\mu(\theta_\alpha)$ to a constant:

$$\mu(\theta_\alpha) = \mu \quad (8)$$

and $k(\theta_\alpha, \theta'_\alpha)$ to the RBF kernel:

$$k^{\text{RBF}}(\theta_\alpha, \theta'_\alpha) = \sigma^2 \prod_\alpha \exp\left(-\frac{(\theta_\alpha - \theta'_\alpha)^2}{2\ell_\alpha^2}\right) \quad (9)$$

where μ, σ, ℓ_α are hyperparameters that can be fixed with maximum likelihood estimation of type II (MLE-II) [18].

In most applications, the energy $E(\theta_\alpha)$ of Eq. (5) is not only C^∞ , but 2π -periodic for each $\alpha = 1, \dots, d$. This property can be imposed to a GP using as covariance function the periodic kernel [19]:

$$k^P(\theta_\alpha, \theta'_\alpha) = \sigma^2 \prod_\alpha \exp\left(-\frac{2}{\ell_\alpha^2} \sin^2\left(\frac{\theta_\alpha - \theta'_\alpha}{2}\right)\right) \quad (10)$$

4.2 Acquisition function

The circuit parameter values of each optimization step are chosen to be the maximum of an appropriately defined acquisition function. A common choice is the expected improvement (EI) acquisition function [20]. Calling $E_{\min} \equiv \min(E_{1\alpha}, \dots, E_{n\alpha})$ the minimum of the previously measured energies and $\hat{E}(\theta_\alpha)$ the energy prediction given by the surrogate model, the EI is defined as:

$$a_{\text{EI}}(\theta_\alpha) \equiv \mathbb{E}_{\hat{E}(\theta_\alpha)}[\max(0, E_{\min} - \hat{E}(\theta_\alpha))] \quad (11)$$

where the expectation value $\mathbb{E}_{\hat{E}(\theta_\alpha)}[\dots]$ is evaluated among all the possible values of $\hat{E}(\theta_\alpha)$ evaluated from the surrogate model.

A great feature of the EI is that both $a_{\text{EI}}(\theta_\alpha)$ and its gradient $\partial_\alpha a_{\text{EI}}(\theta_\alpha)$ are available in closed form if a GP is used as surrogate model [11]. While the EI proves to be very effective for noiseless objective functions, it is not as effective in presence of statistical noise [21]. Since this is the case for VQE, a better choice is an extension of the EI called noisy expected improvement (NEI) [22].

The NEI is an extension of the EI that is not available in closed form. However, it is possible to efficiently evaluate it with (quasi-)Monte Carlo methods. Let $\mathcal{E}_1, \dots, \mathcal{E}_K$ be noiseless (quasi)random

energy functions sampled from the posterior of Eq. (7), and let $a_{\text{EI}}(\theta_\alpha|\mathcal{E}_1), \dots, a_{\text{EI}}(\theta_\alpha|\mathcal{E}_K)$ be EIs defined over them. Then:

$$a_{\text{NEI}}(\theta_\alpha) \simeq \frac{1}{K} \sum_i a_{\text{EI}}(\theta_\alpha|\mathcal{E}_i) \quad (12)$$

Each summand of Eq. (12) is available in closed form so that its evaluation can be done fast. Furthermore, also its gradient is computable analytically, which can be useful for optimizing the acquisition function.

5. Outline of the algorithm

Here we describe our proposed algorithm step by step and specify details of its implementation. The Hamiltonian and the parametrized circuit of Eq. (4) are problem dependent. Quantum computing libraries have procedures for evaluating the parametrized energy of Eq. (5) with any value of the d -dimensional circuit parameter θ_α using the estimator in Eq. (2) and its error of Eq. (3). This measurement is obtained performing S shots. We used the Qiskit [23] quantum computing library for our tests.

Once the routine for measuring $E(\theta_\alpha)$ is defined, the Bayesian optimization procedure is entirely implemented on a classical computer. We built our test on top of the libraries Ax [24], BoTorch [25] and GPyTorch [26]. The Bayesian optimization is then performed as follows:

1. Generate n quasi-random d dimensional points $\theta_{1\alpha}, \dots, \theta_{n\alpha} \in [0, 2\pi]^d$ with a Sobol sequence [27]. In our tests, we used $n = 3$.
2. Given $\theta_{1\alpha}, \dots, \theta_{n\alpha}$, measure their corresponding energies E_1, \dots, E_n and their errors $\Delta E_1, \dots, \Delta E_n$ as described in Eq. (2) and Eq. (3).
3. Use MLE-II to infer the prior hyperparameter μ of Eq. (8) and σ, ℓ_α of Eq. (9) or Eq. (10), depending on whether RBF or periodic kernel was chosen. The default settings of Ax, BoTorch and GPyTorch were used for this inference.
4. Compute the GP posterior mean and covariance of Eq. (7).
5. The current estimation of the parameters for the global minimum point θ_α^{\min} is chosen among $\theta_{1\alpha}, \dots, \theta_{n\alpha}$ as the $\theta_{i\alpha}$ with the minimum expected energy according to the GP model found at point 4:

$$\theta_\alpha^{\min} \equiv \operatorname{argmin}_{\theta_{i\alpha}} \mu(\theta_{i\alpha}|E_i) \quad (13)$$

and the corresponding estimation of the minimum energy E^{\min} is:

$$E^{\min} \equiv \mu(\theta_\alpha^{\min}|E_i) \quad (14)$$

6. Sample K noiseless energy functions from the posterior GP found at point 4. For each of the K samples, perform a different noiseless GPR using the same hyperparameters found in point 3, and compute the K EIs necessary for the NEI approximation of Eq. (12). In our tests, we used $K = 20$.

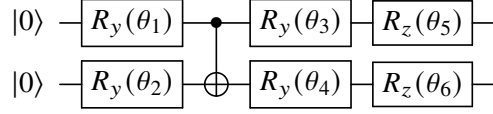
7. Perform a global optimization of the approximated NEI to find its maximum point $\theta_\alpha^{\text{NEI}}$. We performed this optimization with the default procedure of BoTorch, which is a multistart L-BFGS-B [28], where 20 restart points are selected as those with the maximum acquisition function value out of 1000 points drawn from the Sobol sequence in $[0, 2\pi]^d$. The SciPy [29] implementation of L-BFGS-B is used for this step.
8. Add the NEI maximum found at point 7 to the parameter set $\theta_{n+1, \alpha} \leftarrow \theta_\alpha^{\text{NEI}}$ and iterate from point 2 with $n \leftarrow n + 1$ until a break condition is reached which could be realized, e.g., when the minimum of Eq. (13) is stable for a certain number of iterations. Alternatively a fixed number N of iterations might be chosen beforehand in order to keep the number of quantum measurements under control.

6. Testing with IBMQ

We tested the algorithm described in section 5 on a simple two qubits Hamiltonian of the transverse-field Ising model with coupling set to one:

$$\mathcal{H} = -\sigma_x^1 \otimes \sigma_x^2 - \sigma_z^1 - \sigma_z^2 \quad (15)$$

The parametrization of Eq. (4) was achieved using the following quantum circuit:



This circuit was constructed using the procedure described in [30]. It does not have redundant parameters and can cover the whole Hilbert space if we exclude states which are equivalent after the application of a global phase.

For the quantum measurement, we used the Qiskit simulator using the noise model of IBMQ Santiago quantum device. Assuming that only a fixed total number of shots is at our disposal, we tried different algorithms with different number of shots per measurement in order to find the setup that uses this assumed budget most efficiently. We first analyze the results obtained with two algorithms that have been proposed for this specific task: the SPSA [16] in its implementation [12] available in Qiskit, and the NFT [15]². Then we compare their performance with two implementations of BO using respectively the RBF and the periodic kernels of Eqs. (9) and (10). Each algorithm is tested using 20, 40 and 80 number of measurement, each respectively obtained with 64, 32 and 16 number of shots. Therefore, the total number of shots used is 1280 in all cases.

The optimization results are then compared to the exact values of the ground state and the ground state energy. In particular, it is possible to evaluate the state fidelity of a parametrized state with respect to the exact ground state. The fidelity is equal to the square of the scalar product of two states and it quantifies their proximity as its maximum value of 1 is reached when the two states are identical.

After each iteration of the tested algorithms, both the energy and the fidelity have been recorded. The results shown are the average of what was obtained with 20 independent runs starting with

²Its implementation is available at <https://github.com/ken-nakanishi/nftopt>.

different random initial conditions. Their error corresponds to the standard error computed out of these 20 repetitions.

SPSA The results for SPSA³ are shown in Fig. 1. The performances slightly improve decreasing the number of shots per measurements, but none of the three setups gets close to the exact solution.

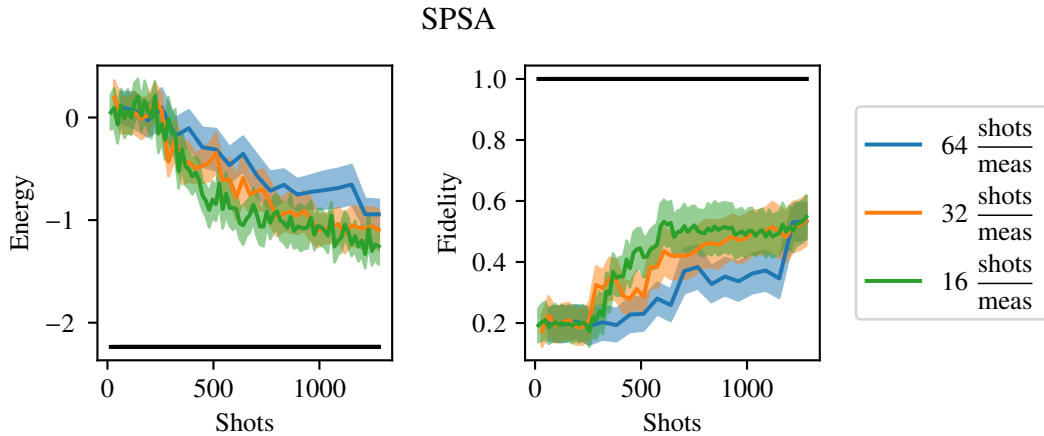


Figure 1: shots/meas indicate the number of shots used for each energy measurement. Black lines correspond to the exact solution.

NFT The performances of the NFT algorithm⁴ are shown in Fig. 2. In this case, the optimization gets quickly close to the solution in all three cases. The speed of convergence increases reducing the number of shots per measurement, but, on the other hand, the stability and the precision of the solution decreases. The fidelity is slightly unstable only with 16 shots per measurement, while energy measurements do not have a good precision with 16 and 32 shots.

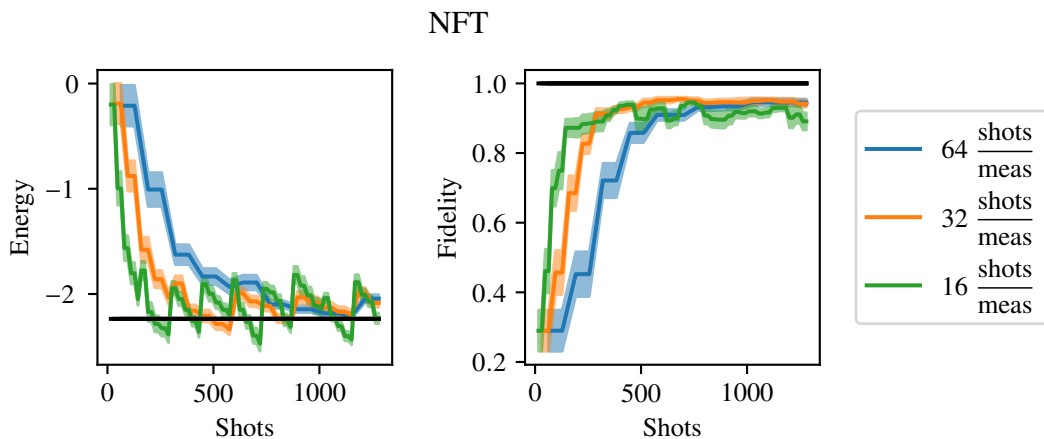


Figure 2: Here is used the same notation of Fig. 1

The NFT algorithm has overall a very quick convergence rate and a low requirement of CPU resources due to its usage of an analytical formula of the target energy. However, its solutions do

³The Qiskit implementation of SPSA was used with its default settings.

⁴These results were obtained setting the variable `reset_interval=4` in the function made available by the authors.

not have a good precision with very small number of shots per measurements as this algorithm does not have built-in methods to infer the real value of the energy, but it has to rely on the mean value of Eq. (2).

Bayesian optimizer Results obtained with the RBF kernel are reported in Fig. 3, while in Fig. 4 we show those obtained with the periodic kernel. The two implementations behave in a very similar way. In both cases the performances improve reducing the number of shots per measurement and the solution gets closer to the exact values without a loss of stability.

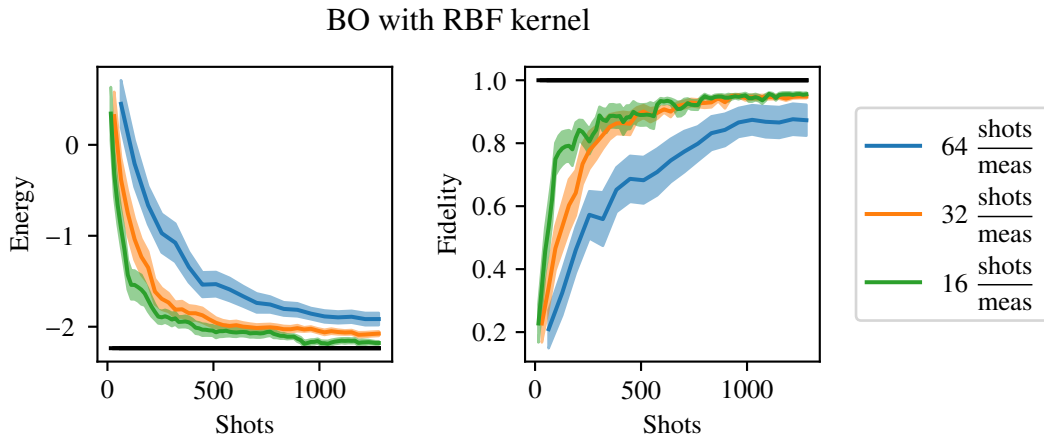


Figure 3: Here is used the same notation of Fig. 1

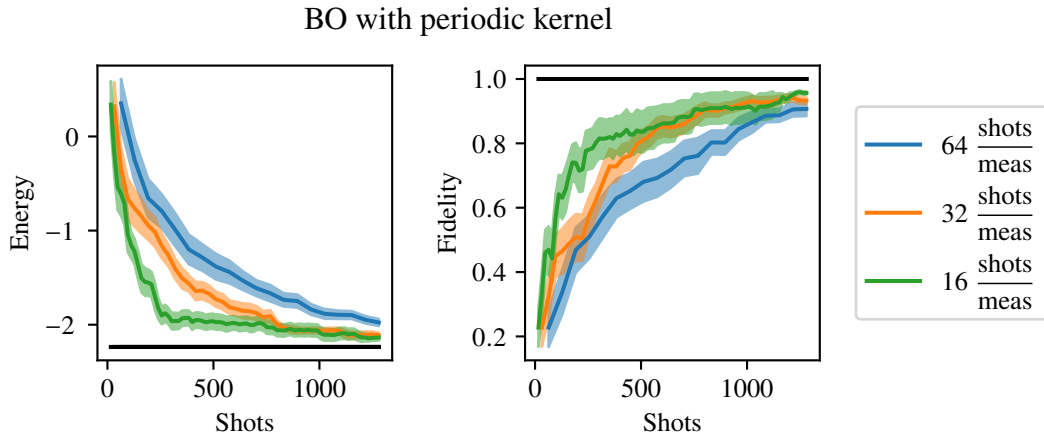


Figure 4: Here is used the same notation of Fig. 1

Comparison of the algorithms In Fig. 5 we report the results of each of the considered algorithms in their best setup, which is 32 shots per measure for NFT and 16 for the others. SPSA is outperformed by NFT and the BOs, which have similar performances in this setup. However, the BOs have better estimation of the ground state energy as its value is inferred with GPR. It removes the tradeoff present in NFT between precision and speed at the cost of increasing CPU time. RBF initially converges faster than periodic, but at the end the solution with the highest fidelity was found with the periodic kernel.

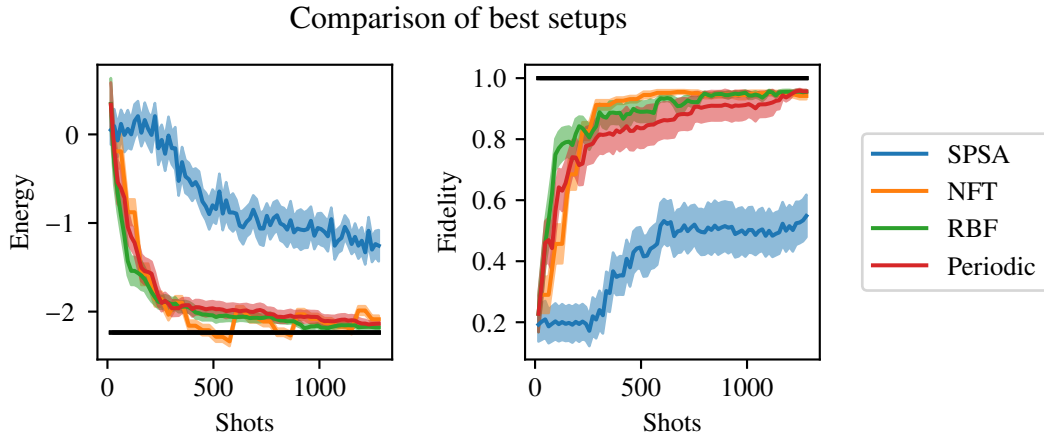


Figure 5: Here are compared the best results of each previous plot.

7. Conclusions and outlooks

Our conclusions are that BO using GPR and NEI is a good choice for VQE in case of noisy measurements obtained with a few number of shots as it can use both the energy mean value of Eq. (2) and its error of Eq. (3). It outperforms SPSA in the here considered case and has a convergence rate similar to the one obtained with NFT. However, BO provides a more precise estimation of the ground state energy, although at the cost of increasing the CPU time, which is presently clearly not a bottleneck when compared to the available QPU time. BO with RBF kernel started converging faster than with periodic kernel, but the most accurate solution was found with the periodic kernel.

The main weakness of BO is its expansion to a high number of circuit parameters, at its standard implementation would require too much CPU time and memory. With some modifications, BO was, however, successfully used with a high number of parameters in other contexts (see for example [31–33]). At the moment we are exploring different possible ways to expand the BO of VQEs to cases in which the number of circuit parameters grows up to $O(100)$, which would provide the BO approach a promising perspective to be used in VQE also for the next generation of quantum computers.

8. Acknowledgements

This project has received funding from the Marie Skłodowska-Curie European Joint Doctorate program STIMULATE of the European Commission under grant agreement No. 765048. G.I.’s position was funded under this program.

References

- [1] J. Biamonte, P. Wittek, N. Pancotti, P. Rebentrost, N. Wiebe and S. Lloyd, *Quantum machine learning*, *Nature* **549** (2017) 195.
- [2] B.P. Lanyon, J.D. Whitfield, G.G. Gillett, M.E. Goggin, M.P. Almeida, I. Kassal et al., *Towards quantum chemistry on a quantum computer*, *Nature chemistry* **2** (2010) 106.
- [3] A. Montanaro, *Quantum algorithms: an overview*, *npj Quantum Information* **2** (2016) 1.
- [4] J. Preskill, *Quantum computing in the nisq era and beyond*, *Quantum* **2** (2018) 79.
- [5] P.W. Shor, *Fault-tolerant quantum computation*, in *Proceedings of 37th Conference on Foundations of Computer Science*, pp. 56–65, IEEE, 1996.
- [6] M.C. Banuls, R. Blatt, J. Catani, A. Celi, J.I. Cirac, M. Dalmonte et al., *Simulating lattice gauge theories within quantum technologies*, *The European physical journal D* **74** (2020) 1.
- [7] A. Kan, L. Funcke, S. Kühn, L. Dellantonio, J. Zhang, J.F. Haase et al., *Investigating a 3+ 1d topological θ -term in the hamiltonian formulation of lattice gauge theories for quantum and classical simulations*, *arXiv preprint arXiv:2105.06019* (2021) .
- [8] J. Zhang, R. Ferguson, S. Kühn, J.F. Haase, C. Wilson, K. Jansen et al., *Simulating gauge theories with variational quantum eigensolvers in superconducting microwave cavities*, *arXiv preprint arXiv:2108.08248* (2021) .
- [9] A. Peruzzo, J. McClean, P. Shadbolt, M.-H. Yung, X.-Q. Zhou, P.J. Love et al., *A variational eigenvalue solver on a photonic quantum processor*, *Nature communications* **5** (2014) 4213.
- [10] H.J. Kushner, *A new method of locating the maximum point of an arbitrary multipeak curve in the presence of noise*, .
- [11] D.R. Jones, M. Schonlau and W.J. Welch, *Efficient global optimization of expensive black-box functions*, *Journal of Global optimization* **13** (1998) 455.
- [12] A. Kandala, A. Mezzacapo, K. Temme, M. Takita, M. Brink, J.M. Chow et al., *Hardware-efficient variational quantum eigensolver for small molecules and quantum magnets*, *Nature* **549** (2017) 242.
- [13] L. Funcke, T. Hartung, K. Jansen, S. Kühn, P. Stornati and X. Wang, *Measurement error mitigation in quantum computers through classical bit-flip correction*, *arXiv preprint arXiv:2007.03663* (2020) .
- [14] J.A. Nelder and R. Mead, *A simplex method for function minimization*, *The computer journal* **7** (1965) 308.
- [15] K.M. Nakanishi, K. Fujii and S. Todo, *Sequential minimal optimization for quantum-classical hybrid algorithms*, *Physical Review Research* **2** (2020) 043158.

- [16] J.C. Spall et al., *Multivariate stochastic approximation using a simultaneous perturbation gradient approximation*, *IEEE transactions on automatic control* **37** (1992) 332.
- [17] G. Matheron, *Traité de géostatistique appliquée. 1 (1962)*, vol. 1, Editions Technip (1962).
- [18] C.K. Williams and C.E. Rasmussen, *Gaussian processes for machine learning*, vol. 2, MIT press Cambridge, MA (2006).
- [19] D.J. MacKay et al., *Introduction to gaussian processes*, *NATO ASI series F computer and systems sciences* **168** (1998) 133.
- [20] J. Močkus, *On bayesian methods for seeking the extremum*, in *Optimization techniques IFIP technical conference*, pp. 400–404, Springer, 1975.
- [21] E. Vazquez, J. Villemonteix, M. Sidorkiewicz and E. Walter, *Global optimization based on noisy evaluations: an empirical study of two statistical approaches*, in *Journal of Physics: Conference Series*, vol. 135, p. 012100, IOP Publishing, 2008.
- [22] B. Letham, B. Karrer, G. Ottoni, E. Bakshy et al., *Constrained bayesian optimization with noisy experiments*, *Bayesian Analysis* **14** (2019) 495.
- [23] G. Aleksandrowicz, T. Alexander, P. Barkoutsos, L. Bello, Y. Ben-Haim, D. Bucher et al., *Qiskit: An open-source framework for quantum computing*, *Accessed on: Mar* **16** (2019) .
- [24] E. Bakshy, L. Dworkin, B. Karrer, K. Kashin, B. Letham, A. Murthy et al., “Ae: A domain-agnostic platform for adaptive experimentation.”
- [25] M. Balandat, B. Karrer, D. Jiang, S. Daulton, B. Letham, A.G. Wilson et al., *Botorch: A framework for efficient monte-carlo bayesian optimization*, *Advances in Neural Information Processing Systems (NeurIPS)* (2020) .
- [26] J. Gardner, G. Pleiss, K.Q. Weinberger, D. Bindel and A.G. Wilson, *Gpytorch: Blackbox matrix-matrix gaussian process inference with gpu acceleration*, in *Advances in Neural Information Processing Systems*, pp. 7576–7586, 2018.
- [27] I.M. Sobol’, *On the distribution of points in a cube and the approximate evaluation of integrals*, *Zhurnal Vychislitel’noi Matematiki i Matematicheskoi Fiziki* **7** (1967) 784.
- [28] R.H. Byrd, P. Lu, J. Nocedal and C. Zhu, *A limited memory algorithm for bound constrained optimization*, *SIAM Journal on scientific computing* **16** (1995) 1190.
- [29] P. Virtanen, R. Gommers, T.E. Oliphant, M. Haberland, T. Reddy, D. Cournapeau et al., *Scipy 1.0: fundamental algorithms for scientific computing in python*, *Nature methods* **17** (2020) 261.
- [30] L. Funcke, T. Hartung, K. Jansen, S. Kühn and P. Stornati, *Dimensional expressivity analysis of parametric quantum circuits*, *Quantum* **5** (2021) 422.

- [31] Z. Wang, F. Hutter, M. Zoghi, D. Matheson and N. de Feitas, *Bayesian optimization in a billion dimensions via random embeddings*, *Journal of Artificial Intelligence Research* **55** (2016) 361.
- [32] C. Li, S. Gupta, S. Rana, V. Nguyen, S. Venkatesh and A. Shilton, *High dimensional bayesian optimization using dropout*, *arXiv preprint arXiv:1802.05400* (2018) .
- [33] J. Kirschner, M. Mutnỳ, N. Hiller, R. Ischebeck and A. Krause, *Adaptive and safe bayesian optimization in high dimensions via one-dimensional subspaces*, *arXiv preprint arXiv:1902.03229* (2019) .