

Machine-Learning-Enhanced Optimization of Noise-Resilient Variational Quantum Eigensolvers

Kim A. Nicoli,^{*a,b,**} Luca J. Wagner^{*a,b,**} and Lena Funcke^{*a,b*}

^aTransdisciplinary Research Area "Building Blocks of Matter and Fundamental Interactions" (TRA Matter), University of Bonn, Regina-Pacis-Weg 3, 53113 Bonn, Germany

^bHelmholtz Institute for Radiation and Nuclear Physics (HISKP), University of Bonn, Nussallee 14-16, 53115 Bonn, Germany

E-mail: knicoli@uni-bonn.de, lwagner@hiskp.uni-bonn.de, lfuncke@uni-bonn.de

Variational Quantum Eigensolvers (VQEs) are a powerful class of hybrid quantum-classical algorithms designed to approximate the ground state of a quantum system described by its Hamiltonian. VQEs hold promise for various applications, including lattice field theory. However, the inherent noise of Noisy Intermediate-Scale Quantum (NISQ) devices poses a significant challenge for running VQEs as these algorithms are particularly susceptible to noise, e.g., measurement shot noise and hardware noise.

In a recent work, it was proposed to enhance the classical optimization of VQEs with Gaussian Processes (GPs) and Bayesian Optimization, as these machine-learning techniques are well-suited for handling noisy data. In these proceedings, we provide additional insights into this new algorithm and present further numerical experiments. In particular, we examine the impact of hardware noise and error mitigation on the algorithm's performance. We validate the algorithm using classical simulations of quantum hardware, including hardware noise benchmarks, which have not been considered in previous works. Our numerical experiments demonstrate that GP-enhanced algorithms can outperform state-of-the-art baselines, laying the foundation for future research on deploying these techniques to real quantum hardware and lattice field theory setups.

The 41st International Symposium on Lattice Field Theory (LATTICE2024) 28 July - 3 August 2024 Liverpool, UK

*Speaker

© Copyright owned by the author(s) under the terms of the Creative Commons Attribution-NonCommercial-NoDerivatives 4.0 International License (CC BY-NC-ND 4.0).

1. Introduction

Quantum hardware with several hundreds of qubits can already be harnessed to outperform classical computers on specific tasks (see, e.g., Ref. [1]). These tasks, however, currently have no practical applications and are specifically designed to be challenging for classical computers. As we have entered the Noisy Intermediate-Scale Quantum (NISQ) era of quantum computing, the primary challenge is to develop algorithms for NISQ devices that may demonstrate quantum advantage in tasks of practical relevance, such as quantum chemistry [2] or quantum field theories [3, 4].

Particularly relevant for NISQ devices are hybrid quantum-classical algorithms, such as the Variational Quantum Eigensolver (VQE) [5, 6]. VQEs can approximate the ground state of quantum Hamiltonians by classical variational optimization of a parametrized quantum circuit. The quantum computer measures the energy of the quantum system for a given set of parameters, i.e., the objective of the minimization problem, while a classical optimization routine is used to find a better set of parameters. Efficient exploration of the parameter space is crucial for successful optimization.

Nakanishi et al. [7] demonstrated that, under certain conditions, the functional form of the VQE's objective function can be explicitly derived and utilized to develop a more efficient optimization routine. The resulting Nakanishi-Fuji-Todo (NFT) algorithm [7] uses two measurements to *sequentially* optimize one parameter of the quantum circuit at a time, by fitting the objective in a one-dimensional subspace. Indeed, this constitutes a sequential minimal optimization (SMO) [8], which differs from global optimization protocols such as gradient-based approaches.

Despite its efficiency, the NFT algorithm faces several challenges, primarily due to the noise of current quantum hardware. In particular, the algorithm relies on noisy measurements to identify the global minimum in the one-dimensional subspace, which often hinders convergence and complicates optimization. Machine learning offers a promising approach to addressing the challenge of noisy measurements and enhancing the NFT algorithm with more flexibility. In a recent work [9], Nicoli et al. proposed an optimization algorithm called NFT-with-EMICoRe (EMICoRe for short), which extends the NFT framework by incorporating Gaussian Process Regression (GPR) and Bayesian Optimization (BO) for a data-driven and potentially more noise-resilient SMO scheme. However, this work [9] focused solely on measurement shot noise, neglecting the impact of quantum hardware noise, such as decoherence, depolarization, and crosstalk errors.

In these proceedings, we build on prior works [7,9] and demonstrate that EMICoRe outperforms the NFT baseline in the presence of simulated quantum hardware noise. We start in Sec. 2 by reviewing the theoretical foundations of VQE, Gaussian Processes, Bayesian Optimization. In Sec. 3, we introduce the EMICoRe algorithm. This is followed by numerical experiments in Sec. 4, which involve classical simulations of noisy quantum hardware, comparing the results with and without error mitigation. We conclude and provide an outlook in Sec. 5.

2. Theoretical Foundations

2.1 Variational Quantum Eigensolver (VQE)

Computing the ground state energy of a quantum system described by a Hamiltonian H and the corresponding ground state wave function $|\psi_{GS}\rangle$ is a key challenge in many fields of physics.

The Hamiltonian of an N-qubit system can be represented as

$$H = \sum_{\alpha=1}^{T} h_{\alpha} P_{\alpha}, \tag{1}$$

where $P_{\alpha} \in \{X, Y, Z, I\}^{\otimes N}$ is a tensor product of single-qubit Pauli operators X, Y, Z and the identity operator *I*. The h_{α} are real coefficients, and *T* is the number of individual Pauli terms.

In general, it is challenging to compute the ground state wave function $|\psi_{GS}\rangle$ because it requires diagonalizing the Hamiltonian H, which is a $\mathbb{C}^{2^N \times 2^N}$ -dimensional matrix for an N-qubit system and thus scales exponentially with the number of qubits. An alternative approach to exact diagonalization is to compute a variational approximation of the ground state. This can be achieved, for instance, by using the Variational Quantum Eigensolver (VQE) [5, 6], where a parametric quantum state $|\psi(\theta)\rangle$ is prepared on a quantum computer. This quantum state, which depends on D angular parameters $\theta \in [0, 2\pi)^D$, is obtained through a sequence of $D' (\geq D)$ unitary quantum gate operations $U(\theta) = U_1 \cdots U_{D'}$ acting on an initial state $|\psi_0\rangle$. A subset D of these operations are Pauli rotation gates, $U_{d'(d)}(\theta_d) = R_{P_d}(\theta_d) = \exp(-i\theta_d P_d)$, where $P_d \in \{X, Y, Z\}$ are the Pauli matrices and θ_d are the angular parameters. The full variational state can thus be written as $|\psi(\theta)\rangle = U(\theta)|\psi_0\rangle$. After the preparation of the quantum state $|\psi(\theta)\rangle$ on the quantum computer, the state is measured in the basis of the individual Pauli strings to obtain the corresponding *measured energy* given by

$$\tilde{E}(\theta) = E^*(\theta) + \varepsilon,$$
(2)

where
$$E^*(\theta) = \langle \psi(\theta) | H | \psi(\theta) \rangle = \langle \psi_0 | U(\theta)^{\dagger} H U(\theta) | \psi_0 \rangle.$$
 (3)

The energy $E^*(\theta)$ is the *true energy* that would be obtained if the quantum computer were noisefree. Typically, a quantum computer is subject to both Gaussian noise, arising from the statistical nature of repeated measurements (i.e., *measurement shot noise*), and *hardware noise*, which stems from imperfections in qubits, gates, and measurements [10]. Both types of noise are encoded in the noise term ε in Eq. (2).

As shown in the green box of Fig. 1, the VQE protocol can be divided into the following steps:

- 1. Initialize a random quantum state with random parameters θ .
- 2. Measure the expectation value $\tilde{E}(\theta) = \langle \psi(\theta) | H | \psi(\theta) \rangle + \varepsilon$ on the quantum computer.
- 3. Use a classical non-linear optimizer, such as gradient-based or gradient-free methods, to determine new parameters $\hat{\theta}$ that minimize $\tilde{E}(\theta)$.
- 4. Iterate to step 2 until $\tilde{E}(\theta)$ converges to the ground state energy.

Under mild assumptions for the above described variational circuit, Nakanishi et al. derived an analytical expression of the VQE's objective functional [7],

$$E^{*}(\boldsymbol{\theta}) = \boldsymbol{b}^{\top} \cdot \left[\bigotimes_{d=1}^{D} \begin{pmatrix} \cos\left(\theta_{d}\right) \\ \sin\left(\theta_{d}\right) \\ 1 \end{pmatrix} \right], \quad \forall \boldsymbol{\theta} \in [0, 2\pi)^{D}, \tag{4}$$



Figure 1: Illustration of the VQE workflow. For more details, see the text.

where **b** is a set of arbitrary coefficients and θ_d is a *D*-dimensional vector of angular parameters for the variational quantum circuit. In particular, this functional form leads to the following observation: when keeping all but one parameter constant, i.e.,

$$\boldsymbol{\theta}_{d}^{(t)} = (\theta_1, \dots, \theta_{d-1}, \theta_d, \theta_{d+1}, \dots, \theta_D)^{\mathsf{T}} \quad \text{with} \quad \theta_{i \neq d} = \text{const.},$$
(5)

the form of the objective function for this one-dimensional subspace of the whole optimization problem reduces to [7]

$$E^{*(t)}(\theta_{d}^{(t)}) = \sum_{k=1}^{K} h_{k} \left\langle \psi_{k} \left| U_{d}^{\dagger}(\theta_{d}^{(t)}) P_{k} U_{d}(\theta_{d}^{(t)}) \right| \psi_{k} \right\rangle$$

= $a_{1,d}^{(t)} \cos \left(\theta_{d}^{(t)} - a_{2,d}^{(t)} \right) + a_{3,d}^{(t)},$ (6)

where $a_{\ell,d}^{(t)}$ ($\ell = 1, 2, 3$) denote parameters that are independent of $\theta_d^{(t)}$, while the superscript (t) refers to the optimization step. As a result, measuring three points in the one-dimensional subspace is sufficient to fit a function across the entire domain and analytically determine its minimum.

The NFT optimization algorithm [7] uses this property and sequentially optimizes each parameter in the set θ by requiring (at least) two measurements for each optimization step. The new observed points are *deterministically* chosen along the axis *d*, i.e.,

$$\boldsymbol{\Theta}' = \left(\boldsymbol{\theta}_1', \boldsymbol{\theta}_2'\right) = \left\{ \hat{\boldsymbol{\theta}}^{(t-1)} - \alpha \boldsymbol{e}_d, \hat{\boldsymbol{\theta}}^{(t-1)} + \alpha \boldsymbol{e}_d \right\},$$
(7)

where $\hat{\theta}^{(t-1)}$ denotes the previous best point from step (t-1).¹ This approach is advantageous because it enables analytical optimization of each parameter within its respective one-dimensional subspace, eliminating the need to compute gradients as required by gradient-based methods, such as SPSA [12, 13]. Furthermore, gradient-based approaches may optimize all parameters simultaneously during each gradient update but do not guarantee analytical minimization.

¹The original paper from Nakanishi et al. [7] proposed a shift α of $\frac{\pi}{2}$ for Eq. (7). In the work by Nicoli et al. [9] it was suggested that equidistant points for probing the subspace should be used, i.e., a shift of $\frac{2\pi}{3}$ should be preferred. Later, it was derived by Anders et al. [11] that the offset of $\frac{2\pi}{3}$ provides a constant variance over the whole 1D-subspace. In a private communication with the authors, it was reported that, because of a bug, the actual value used in the experiments of [9] was $\frac{\pi}{3}$ instead of $\frac{2\pi}{3}$ as reported in the paper. Although this did not lead to significant differences in the numerical experiments, in order to be consistent with the prior work, we will stick to the convention of using $\frac{\pi}{3}$ for NFT.

2.2 Gaussian Processes and Bayesian Optimization

The NFT algorithm can find the *exact* analytical minimum within the one-dimensional subspace if and only if the state preparation and measurements on the quantum computer are exact—that is, if an infinite number of shots are taken and no hardware noise is present. In a real-world scenario, due to the unavoidable hardware noise on current NISQ hardware and the inherent statistical uncertainty from repeated measurements (shots), it may be beneficial to enhance existing algorithms with frameworks like Bayesian Optimization (BO) [14] and Gaussian Process Regression (GPR) [15]. These techniques, widely established in machine learning, are often well-suited to deal with noisy data [15]. Let $E^*(\cdot) : X \mapsto \mathbb{R}$ be an unknown (black-box) objective function to be minimized (i.e., the energy in our case). With the aid of a function that suggests promising new points to observe, often referred to as *acquisition function* [14], one iteratively approximates the target objective with a surrogate function $E(\theta) \approx E^*(\theta)$, sampled from the GP. At each optimization step, after new points are measured, those are added to the previous set of measured points, used to update the GP.

A common choice for a surrogate model is a GP regression model with a one-dimensional Gaussian likelihood and a GP prior [9],

$$p(\tilde{E} \mid \boldsymbol{\theta}, E(\cdot)) = \mathcal{N}_1\left(\tilde{E}; E(\boldsymbol{\theta}), \sigma^2\right), \quad p(E(\cdot)) = \operatorname{GP}(E(\cdot); \nu(\cdot), k(\cdot, \cdot)).$$
(8)

Here, σ^2 is the variance of the noise ε defined in Eq. (2), and the prior mean function $v(\cdot)$ is set to 0 in this work. In the case of VQEs, a zero-mean prior is a reasonable choice because the target function $E^*(\theta)$, which one aims to model, is a linear combination of functions centered around zero (up to a potential constant offset stemming from Eq. (4)). The function $k(\cdot, \cdot)$ is the kernel (covariance) function encoding the expected correlation between two measured points. We refer to Section 3.1 for more details on the kernel function we use in these proceedings. The GP is trained on noisy measurements $\tilde{E} = E^*(\theta) + \varepsilon$ that were collected throughout previous optimization steps.

Let $\{\Theta, \tilde{E}\}$ be *N* training samples, where $\Theta = (\theta_1, \dots, \theta_N) \in X^N$ are the angular inputs of the parametrized quantum circuit, and $\tilde{E} = (\tilde{E}_1, \dots, \tilde{E}_N)^\top \in \mathbb{R}^N$ are the energies measured for the corresponding angular parameters $\theta_i \in \Theta$. The posterior of a GP regression model, see Eq. (8), is still a GP, i.e., $p(E(\cdot) | \Theta, \tilde{E}) = \text{GP}(E(\cdot); \mu_{\Theta}(\cdot), s_{\Theta}(\cdot, \cdot))$. Thus, in Ref. [9], it was shown that for arbitrary *M* test points $\Theta' = (\theta'_1, \dots, \theta'_M) \in X^M$, the posterior, or predictive distribution of the function values $E' = (E(\theta'_1), \dots, E(\theta'_M))^\top \in \mathbb{R}^M$, is an *M*-dimensional Gaussian,

$$p\left(\boldsymbol{E}' \mid \boldsymbol{\Theta}, \tilde{\boldsymbol{E}}\right) = \mathcal{N}_{M}\left(\boldsymbol{E}'; \boldsymbol{\mu}_{\boldsymbol{\Theta}}', \boldsymbol{S}_{\boldsymbol{\Theta}}'\right),\tag{9}$$

with mean and covariance analytically given by

$$\boldsymbol{\mu}_{\boldsymbol{\Theta}}^{\prime} = \boldsymbol{K}^{\prime \top} \left(\boldsymbol{K} + \sigma^{2} \boldsymbol{I}_{N} \right)^{-1} \boldsymbol{\tilde{E}}, \quad \boldsymbol{S}_{\boldsymbol{\Theta}}^{\prime} = \boldsymbol{K}^{\prime \prime} - \boldsymbol{K}^{\prime \top} \left(\boldsymbol{K} + \sigma^{2} \boldsymbol{I}_{N} \right)^{-1} \boldsymbol{K}^{\prime}.$$
(10)

The notation with a prime refers to new, unseen points with the corresponding energies predicted by the GP. Here, $\mathbf{K} = k(\mathbf{\Theta}, \mathbf{\Theta}) \in \mathbb{R}^{N \times N}$, $\mathbf{K}' = k(\mathbf{\Theta}, \mathbf{\Theta}') \in \mathbb{R}^{N \times M}$, and $\mathbf{K}'' = k(\mathbf{\Theta}', \mathbf{\Theta}') \in \mathbb{R}^{M \times M}$ are the train, train-test, and test kernel matrices, respectively, where $k(\mathbf{\Theta}, \mathbf{\Theta}')$ denotes the kernel matrix evaluated at each column of $\mathbf{\Theta}$ and $\mathbf{\Theta}'$, such that $(k(\mathbf{\Theta}, \mathbf{\Theta}'))_{n,m} = k(\theta_n, \theta'_m)$. Moreover, $I_N \in \mathbb{R}^{N \times N}$ denotes the identity matrix. The subscript $\mathbf{\Theta}$ in Eq. (9) denotes which input points were used to train the GP. Note that the predictive mean and covariance are given in their vector and matrix form, respectively. The predictive mean can be computed as a function of one arbitrary test point θ' , i.e., $\mu_{\Theta}(\theta') = K'^{\top} (K + \sigma^2 I_N)^{-1} \tilde{E}$ with $K' = k (\Theta, \theta') \in \mathbb{R}^N$. The computation of a single predictive covariance $s_{\Theta}(\cdot, \cdot)$ with $S'_{\Theta} = s_{\Theta}(\Theta', \Theta') \in \mathbb{R}^{M \times M}$ follows analogously. For more details on how to compute the predictive mean and variance, we refer to Ref. [9].

As the number of measurements increases, the knowledge gained by the GP during optimization can be used not only for regression but also to choose the next points to observe. This is achieved by combining GPR and BO and solving the following maximization problem:

$$\max_{\mathbf{\Theta}'} a_{\mathbf{\Theta}^{(t-1)}}(\mathbf{\Theta}')$$

where $a_{\Theta^{(t-1)}}(\cdot)$ is the so-called *acquisition function*. This function uses a GP model trained on previously observed data pairs $\{\Theta^{(t-1)}, \tilde{E}\}$, collected over the prior (t-1) iterations, to assess how promising a set of new input points Θ' may be. The points that maximize the acquisition function are likely to be the most informative for the next iteration.

3. Method

In Section 3.1, we leverage the functional form of the VQE objective [7] to construct a *physics-informed kernel* suited for GPR [9]. In Section 3.2, we introduce a novel type of acquisition function called *Expected Maximum Improvement over Confident Regions* (EMICoRe) [9]. In Section 3.2.1, we combine these concepts and present the *NFT-with-EMICoRe algorithm*, called *EMICoRe algorithm* for short [9]. This algorithm demonstrates how the NFT baseline algorithm can be improved to achieve a more flexible and noise-resilient algorithm when combined with BO and GPR.

3.1 VQE Kernel

A GP needs to be equipped with a kernel (covariance) function. While there are many types of such functions, one popular example is the Gaussian-RBF kernel [15]. This kernel encodes the assumption that input vectors that are close in the input space exhibit stronger correlations compared to those that are distant, with the correlation strength decaying exponentially as a function of the distance. While this may be a reasonable assumption, an exponential decay in correlation, as seen in the Gaussian-RBF kernel, may not accurately reflect the underlying features of every problem. For this reason, the choice of kernel in a GPR task can significantly affect performance. In Ref. [9] the authors proposed the so-called VQE kernel:

$$k^{\text{VQE}}\left(\boldsymbol{\theta},\boldsymbol{\theta}'\right) = \sigma_0^2 \prod_{d=1}^{D} \left(\frac{\gamma^2 + \cos\left(\theta_d - \theta_d'\right)}{1 + \gamma^2}\right),\tag{11}$$

where *D* is the number of angular parameters of the parametrized quantum circuit, σ_0 corresponds to the prior variance, and $\gamma^2 \ge 1$ controls the smoothness of the kernel. Both σ_0 and γ^2 are hyperparameters of the kernel.

In [9], it was shown that the VQE kernel can be decomposed as

$$k^{\text{VQE}}\left(\boldsymbol{\theta},\boldsymbol{\theta}'\right) = \boldsymbol{\phi}(\boldsymbol{\theta})^{\top}\boldsymbol{\phi}(\boldsymbol{\theta}'), \quad \text{with} \quad \boldsymbol{\phi}(\boldsymbol{\theta}) = \sigma_0 \left(1 + \gamma^2\right)^{-D/2} \left[\bigotimes_{d=1}^{D} \left(\gamma, \cos\theta_d, \sin\theta_d\right)^{\top} \right], \quad (12)$$

where the feature maps $\phi(\theta)$ fulfill the same VQE functional form as derived in Eq. (4). Since a GP regression model is only defined in terms of inner products of the feature maps, one just requires access to the kernel form without explicitly specifying ϕ —a procedure known as *kernel trick* [15].

By using the VQE kernel instead of, for example, the Gaussian-RBF kernel, the GP is ensured to sample functions that align with the required functional form from Eq. (4). Consequently, as described in Eq. (6), when optimizing over a one-dimensional subspace—i.e., varying only one parameter at a time—three noiseless observed points are sufficient to uniquely identify the target function [7] across the entire subspace. However, in the presence of measurement shot noise, the target function in the subspace can only be identified up to some statistical uncertainty. Therefore, the GP provides not only energy estimates but also their corresponding uncertainties for new measurements in the subspace.

Furthermore, it has been shown that the decomposition of the VQE kernel in Eq. (12) can be generalized so that each entry of θ may parametrize multiple gates at once [9]. This property becomes crucial if the Hamiltonian is invariant under specific symmetries, leading to multiple parameters being identical and requiring simultaneous updates.

3.2 Expected Maximum Improvement over Confident Regions

To perform BO in combination with GPR, one should pick a suitable acquisition function to identify the most promising points to measure at each optimization step. A novel acquisition function termed *Expected Maximum Improvement over Confident Regions* (EMICoRe) was introduced in the context of GPR for VQEs [9]. This acquisition function relies on the idea of *Confident Region* (CoRe) mathematically defined as

$$\mathcal{Z}_{\Theta} = \left\{ \boldsymbol{\theta} \in \mathcal{X}; s_{\Theta}(\boldsymbol{\theta}, \boldsymbol{\theta}) \le \kappa^2 \right\}.$$
(13)

The CoRe describes all points for which the *predictive* covariance $s_{\Theta}(\cdot, \cdot)$ is smaller than or equal to a threshold κ^2 , which is a hyperparameter of the model. For sufficiently small κ , an appropriate kernel function, and a weak prior—i.e., enough previously observed points—every point in the CoRe can be regarded as "already observed" due to the sufficiently small uncertainty of the GP. Leveraging the concept of CoRe, Nicoli et al. [9] proposed the following acquisition function:

$$a^{\text{EMICoRe}}\left(\mathbf{\Theta}'\right) = \frac{1}{M} \left\langle \max\left(0, \min_{\theta \in \mathcal{Z}_{\Theta}} E(\theta) - \min_{\theta \in \mathcal{Z}_{\widetilde{\Theta}}} E(\theta)\right) \right\rangle_{p(E(\cdot)|\mathbf{\Theta}, \widetilde{E})},$$
(14)

where $\widetilde{\Theta} = (\Theta, \Theta') \in \mathcal{X}^{N+M}$ denotes the augmented training set with the new input points $\Theta' \in \mathcal{X}^M$ and the previously observed ones $\Theta \in \mathcal{X}^N$.

The term $\min_{\theta \in \mathcal{Z}_{\Theta}} E(\theta)$ in Eq. (14) corresponds to the estimated minimum of the CoRe, identified by all previously observed points Θ . This value is drawn from the GP's posterior distribution and identifies the estimated minimum energy for any choice of circuit parameters previously observed. The other term in Eq. (14), $\min_{\theta \in \mathcal{Z}_{\widetilde{\Theta}}} E(\theta)$, is the estimated minimum over the augmented training set $\widetilde{\Theta} = (\Theta, \Theta')$. Note that within this augmented set, the new input points $\Theta' \in \mathcal{X}^M$ can be treated as if they had already been observed. This is because, if a new input point lies within the CoRe, the GP's uncertainty is below a given threshold, allowing the reasonable assumption that the GP is *highly confident* close to that point. Consequently, this point can be implicitly considered *known*, i.e., as if it has been observed, despite no measurement being performed on the quantum computer.

3.2.1 The EMICoRe algorithm

To address the challenges posed by noisy measurements on quantum hardware, Nicoli et al. [9] proposed to combine NFT with GP regression and BO, thus merging all the tools previously introduced in Sections 2.2, 3.1 and 3.2.

The proposed EMICoRe algorithm builds upon the structure of the NFT algorithm, incorporating two key modifications:

- 1. Learnable shifts: Instead of taking a fixed shift α to identify two new (equidistant) candidate points, see Eq. (7), the EMICoRe algorithm uses the EMICoRe acquisition function from Section 3.2 to identify the most *informative* points to measure next.
- 2. **Re-use of previous measurements**: At each iteration, the minimum on the one-dimensional subspace is estimated using Least Square Regression with the posterior mean of the GP. This implies that as more measurements are made on the quantum computer and incorporated into the surrogate model, the GP will provide increasingly accurate estimates.

These modifications distinguish the proposed algorithm from NFT, where the minimum estimation on the subspace depends solely on the previous best point and the *fixed* shifts for the new measurements. Both modifications aim to make EMICoRe a more flexible and resilient algorithm in the presence of quantum noise. The EMICoRe algorithm starts with T_{NFT} steps of the baseline NFT algorithm to gather sufficient measurements for instantiating an informative GP that can make predictions. After this initial phase, the EMICoRe algorithm begins learning the shifts at each step, using the previous measurements and storing the new measurements for future predictions. For a detailed description of the EMICoRe algorithm and its subroutine, we refer to the original paper [9].

Figure 2 visualizes the algorithmic procedure of the full EMICoRe algorithm. The objective function $E^*(\theta)$ (true energy) is represented by the red lines, while the red shaded area indicates the standard deviation of a measurement on the quantum computer. The blue lines indicate the predictive mean of the GP $\mu_{\Theta}(\cdot)$, while the shaded area denotes its predictive uncertainty $\sqrt{s_{\Theta}(\cdot, \cdot)}$, i.e., one standard deviation. The figure should be read from top left to bottom right following the black arrow.

The top left plot shows the optimization in the direction θ_d for $d \in \{1, \ldots, D\}^2$. The quantity $\hat{\theta}^{(t)}$ denotes the set of parameters which give the lowest energy $E(\hat{\theta}^{(t)})$ from the GP after the *t*-th optimization step. Note that at this point, the GP has access to a dataset of previous measurements $\{\Theta^{(t)}, \tilde{E}^{(t)}\}$. This explains why the blue line, along with the corresponding uncertainty, provides a good prediction of the red curve along direction *d*, even before any optimization has been performed.

The EMICoRe subroutine is now responsible for finding the promising points to be measured next, i.e., Θ' in Eq. (7) with the shifts α being different for θ'_1 and θ'_2 . First, the subroutine divides the subspace into a grid of equally spaced points, as represented by the grid lines at the bottom of the second plot. Second, the subroutine prepares all possible combinations of candidate pairs $\Theta'_k = \{\theta'_{k,2}, \theta'_{k,2}\}$. These points are added to the *base-GP* in pairs, allowing a discretized CoRe to be computed for each pair (orange vertical lines). The subroutine then computes the acquisition function for each candidate pair Θ'_k using the respective CoRe. The pair that maximizes the

²Note the direction to be optimized is chosen sequentially until d = D, after which the algorithm restarts from d = 0.



Figure 2: Visualization of the EMICoRe algorithm. We refer to the main text for more details.

acquisition function, as described in Eq. (14), is returned as Θ' . According to the algorithm, this pair represents the points that are the most promising to be observed next. The fifth plot displays the measurements on the quantum computer at the suggested points, indicated by the blue crosses. The GP is finally updated using the new measurements—note the change of the blue curve and the shaded area in the plot. Finally, the GP's posterior mean is computed for three equidistant points on the one-dimensional subspace. These points are then used to fit a cosine function, according to the functional form of the subspace given by Eq. (6). This fitted cosine function is then used to obtain the analytic minimum of the one-dimensional subspace.

We recall that, due to the special form of the VQE-kernel and the design of the EMICoRe algorithm, the function sampled from the GP—i.e, the posterior mean—follows the same functional form as in Eq. (4). Therefore, fitting a sinusoidal function to the GP's posterior is a valid assumption.

4. Numerical Experiments

In the following, we perform numerical experiments to demonstrate the performance of EMI-CoRe in the presence of simulated hardware noise. Our experiments focus on the quantum Heisenberg Hamiltonian, a typical Hamiltonian for benchmarking the performance of VQEs,

$$H = -\left[\sum_{j=1}^{Q-1} \left(J_X X_j X_{j+1} + J_Y Y_j Y_{j+1} + J_Z Z_j Z_{j+1}\right) + \sum_{j=1}^{Q} \left(h_X X_j + h_Y Y_j + h_Z Z_j\right)\right], \quad (15)$$

where the index of the Pauli matrices $\{X_j, Y_j, Z_j\}$ refers to the qubit that the Pauli gate is applied to, and open boundary conditions are considered. In this work, we choose the couplings to be *at the critical point* with $(J_X, J_Y, J_Z) = (-1, 0, 0)$ and $(h_X, h_Y, h_Z) = (0, 0, -1)$. The resulting Hamiltonian is also known as the Ising Hamiltonian *at criticality*.

In order to assess the performance of the VQE during optimization, we mainly consider two metrics: *energy* and *fidelity*. The former is the lowest energy reached by the VQE during the optimization process, i.e., the expectation value of the Hamiltonian with respect to the variational state. The latter measures the overlap between the variational state and the true ground state obtained with exact diagonalization. We note that the true ground-state energy $\underline{E}^* = \langle \psi_{GS} | H | \psi_{GS} \rangle$, along with with the corresponding true ground-state wave function $|\psi_{GS}\rangle$, can be obtained in our case, as the size of the systems we address is sufficiently small. This is important as it provides access to a ground truth for benchmarking different algorithms, such as NFT versus EMICoRe. We refer to [9] for more details on the metrics we used. To test the stability of the optimization with respect to initialization, we initialize several independent runs with different seeds. Specifically, we use 50 runs for the experiments without hardware noise and 10 runs for the experiments with hardware noise, see Section 4.1 for more details. In all plots, we report the median (50th percentile) as a solid line, with the 25th and 75th percentiles as the lower and upper bounds, respectively, represented by shaded regions around the solid line. The results are plotted as a function of the total number of measurements performed on the quantum computer.

For all simulations, we use an *L*-layered and fully connected Efficient SU(2) circuit with $D = 2(L + 1) \cdot Q$ total angular parameters. The classical simulations of the quantum hardware are performed with the Qiskit [16] library, while the Python implementation of the EMICoRe algorithm has been adapted from the original code available on GitHub [17].

All the parameters for EMICoRe are identical to those chosen for the simulation in Fig. 7 of the original paper [9]. As noted in Footnote 1, we used $\alpha = \pi/3$ as the shift for NFT. Moreover, in contrast to the original NFT paper [7], we did not include a stabilization step, as our results indicated that its use often led to poorer performance. Further investigation of this phenomenon is left to future work.

4.1 Hardware Noise Analysis of EMICoRe

Previous work [9] analyzed EMICoRe's performance in the presence of measurement shot noise, showing that EMICoRe outperforms NFT more significantly at higher noise levels, i.e., with fewer shots. In these proceedings, we extend the analysis to include hardware noise as well.

In the following experiments, the setup remains constant: we consider a 5-qubit system with 3-layers of the EfficientSU2 circuit and the Ising Hamiltonian with open boundary conditions. All simulations are conducted with a total of 600 measurements ³, using N = 1024 shots per measurement.

In Fig. 3a, we conduct experiments with simulated hardware noise and no error mitigation. In Fig. 3b and Fig. 3c, we apply *twisted readout error extinction* (TREX) [18] and *zero noise extrapolation* (ZNE) [19, 20] as error mitigation strategies, respectively. Hardware noise is simulated

³Note that two measurements are performed at each step, meaning that 600 measurements correspond to 300 optimization steps.



(c) Classical simulation of noisy quantum hardware with ZNE error mitigation.

Figure 3: Energy (left) and fidelity (right) for NFT (green) and EMICoRe (red). We show results for the critical Ising Hamiltonian using Q = 5 qubits and L = 3 layers, for an Efficient SU(2) quantum circuit, and $N_{\text{shots}} = 1024$ measurement shots. The top row displays runs with simulated hardware noise and no error mitigation, while the second and third rows present results with TREX and ZNE error mitigation schemes, respectively. The median (solid line) and the 25th and 75th percentiles (colored regions surrounding the median) are obtained with 50 independent seeded runs for experiments in Fig. 3a and with 10 runs for Figs. 3b and 3c. A density of the trial distribution at the end of the optimization is shown right next to each plot.

with the Fake5QV1 backend. The results in Fig. 3a are obtained by aggregating 50 independent runs, whereas Figs. 3b and 3c include only 10 runs, as the feature was deprecated before further data could be collected. Table 1 summarizes the results from Fig. 3, reporting the mean and standard deviation computed over multiple seeded trials for simulations with and without error mitigation.

When hardware noise is taken into account, EMICoRe significantly outperforms NFT, as shown in Fig. 3a, demonstrating faster and more reliable convergence. For instance, the trial density on the right-hand-side is more peaked around the median. These empirical results indicate that NFT is considerably more affected by hardware noise compared to EMICoRe. While neither algorithm reaches the ground-state energy, EMICoRe achieves a fidelity of nearly 80%. With error

Table 1: Mean and standard deviation of the energy and fidelity at the final optimization step for all simulations presented in Fig. 3. For the energy, lower values are better; for the fidelity, higher values are better.

	EMICoRe		NFT	
	Energy $E^*(\hat{\theta})$	Fidelity	Energy $E^*(\hat{\theta})$	Fidelity
noise	(-5.472 ± 0.119)	(0.782 ± 0.04)	(-4.972 ± 0.821)	(0.641 ± 0.214)
noise, TREX	(-5.624 ± 0.122)	(0.854 ± 0.051)	(-5.173 ± 0.543)	(0.622 ± 0.245)
noise, ZNE	(-5.611 ± 0.059)	(0.814 ± 0.037)	(-5.03 ± 0.841)	(0.653 ± 0.226)

mitigation, EMICoRe outperforms NFT in both energy minimization and fidelity, demonstrating greater stability across different initializations and achieving lower energies as well as higher fidelities.

5. Conclusion

In these proceedings, we presented a novel machine-learning-enhanced optimization approach for VQE algorithms to find the ground state of the quantum Ising Hamiltonian in the presence of simulated hardware noise. This algorithm, called EMICoRe [9], combines Gaussian Process Regression for noise-resilient energy estimation and Bayesian Optimization to predict new candidate points for measurement at each step of the optimization.

Our numerical experiments demonstrated that EMICoRe outperforms the state-of-the-art baseline, the NFT algorithm [7], in the presence of simulated hardware noise. EMICoRe converges to lower energies than NFT and shows greater robustness across different initializations. This robustness underscores the potential of the proposed method for future applications to more complex quantum systems and under the challenging conditions of real NISQ devices.

Acknowledgments

The authors thank Christopher J. Anders and Shinichi Nakajima for insightful discussions. This project was supported by the Deutsche Forschungsgemeinschaft (DFG, German Research Foundation) as part of the CRC 1639 NuMeriQS – project no. 511713970.

References

- A. Morvan et al., *Phase transitions in random circuit sampling*, *Nature* 634 (2024) 328 [2304.11119].
- [2] A.J. McCaskey, Z.P. Parks, J. Jakowski, S.V. Moore, T.D. Morris, T.S. Humble et al., Quantum chemistry as a benchmark for near-term quantum computers, npj Quantum Information 5 (2019) 99 [1905.01534].
- [3] A. Di Meglio et al., *Quantum computing for high-energy physics: State of the art and challenges*, *PRX Quantum* **5** (2024) 037001 [2307.03236].

- [4] L. Funcke, T. Hartung, K. Jansen and S. Kühn, *Review on quantum computing for lattice field theory*, *PoS* LATTICE2022 (2023) 228 [2302.00467].
- [5] 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 Commun. 5 (2014) 4213
 [1304.3061].
- [6] J.R. McClean, J. Romero, R. Babbush and A. Aspuru-Guzik, *The theory of variational hybrid quantum-classical algorithms*, *New J. Phys.* 18 (2016) 023023 [1509.04279].
- [7] K.M. Nakanishi, K. Fujii and S. Todo, *Sequential minimal optimization for quantum-classical hybrid algorithms*, *Phys. Rev. Res.* **2** (2020) 043158 [1903.12166].
- [8] C.P. John, Sequential minimal optimization: a fast algorithm for training support vector machines, MSRTR: Microsoft Research 3 (1998) 88.
- [9] K.A. Nicoli, C.J. Anders, L. Funcke, T. Hartung, K. Jansen, S. Kühn et al., *Physics-informed bayesian optimization of variational quantum circuits*, *Advances in Neural Information Processing Systems* 36 (2024) 18341 [2406.06150].
- [10] K. Georgopoulos, C. Emary and P. Zuliani, Modeling and simulating the noisy behavior of near-term quantum computers, Phys. Rev. A 104 (2021) 062432 [2102.02109].
- [11] C.J. Anders, K.A. Nicoli, B. Wu, N. Elosegui, S. Pedrielli, L. Funcke et al., Adaptive observation cost control for variational quantum eigensolvers, Proceedings of the 41st International Conference on Machine Learning 235 (2024) 1557 [2502.01704].
- [12] J. Spall, Multivariate stochastic approximation using a simultaneous perturbation gradient approximation, IEEE Transactions on Automatic Control **37** (1992) 332.
- [13] S. Finck and H.-G. Beyer, *Performance analysis of the simultaneous perturbation stochastic approximation algorithm on the noisy sphere model*, *Theor. Comp. Sci.* **419** (2012) 50.
- [14] P.I. Frazier, A tutorial on bayesian optimization, 1807.02811.
- [15] C.E. Rasmussen and C.K.I. Williams, *Gaussian Processes for Machine Learning*, The MIT Press (2005), 10.7551/mitpress/3206.001.0001.
- [16] A. Javadi-Abhari, M. Treinish, K. Krsulich, C.J. Wood, J. Lishman, J. Gacon et al., *Quantum computing with Qiskit*, 2405.08810.
- [17] K.A. Nicoli, C.J. Anders et al., *Emicore: Expected maximum improvement over confident regions, GitHub repository* (2023).
- [18] E. van den Berg, Z.K. Minev and K. Temme, Model-free readout-error mitigation for quantum expectation values, Phys. Rev. A 105 (2022) 032620 [2012.09738].
- [19] K. Temme, S. Bravyi and J.M. Gambetta, Error mitigation for short-depth quantum circuits, Phys. Rev. Lett. 119 (2017) 180509 [1612.02058].

[20] Y. Li and S.C. Benjamin, *Efficient variational quantum simulator incorporating active error minimization*, *Phys. Rev. X* 7 (2017) 021050 [1611.09301].