

Reducing runtime and error in VQE using deeper and noisier quantum circuits

Amara Katarawa¹, Alex Kunitsa¹, Borja Peropadre^{1,*} and Peter Johnson^{1,†}

¹Zapata Computing, Boston, MA 02110 USA

(Dated: October 22, 2021)

The rapid development of noisy intermediate-scale quantum (NISQ) devices has raised the question of whether or not these devices will find commercial use. Unfortunately, a major shortcoming of many proposed NISQ-amenable algorithms, such as the variational quantum eigensolver (VQE), is coming into view: the algorithms require too many independent quantum measurements to solve practical problems in a reasonable amount of time. This motivates the central question of our work: how might we speed up such algorithms in spite of the impact of error on NISQ computations? We demonstrate on quantum hardware that the estimation of expectation values, a core subroutine of many quantum algorithms including VQE, can be improved in terms of precision and accuracy by using a technique we call *robust amplitude estimation*. Consequently, this method reduces the runtime to achieve the same mean-squared error compared to the standard prepare-and-measure estimation method. The surprising result is that by using deeper, and therefore more error-prone, quantum circuits, we realize more accurate quantum computations in less time. As the quality of quantum devices improves, this method will provide a proportional reduction in estimation runtime. This technique may be used to speed up quantum computations into the regime of early fault-tolerant quantum computation and aid in the realization of quantum advantage.

Advances in quantum devices have fueled the development of near-term quantum algorithms [1–4]. A driving question in the field is: can a quantum device achieve quantum advantage using such algorithms well-before the era of fault tolerance? The variational quantum eigensolver (VQE), used for simulating chemistry and materials, has been a promising contender [5]. However, recent work has shown that the standard approach to implementing VQE requires exceedingly long runtimes for problem instances of industrial relevance [6]. This bottleneck is referred to as the “measurement problem”. Beyond quantum simulation, the measurement problem plagues quantum machine learning and near-term quantum algorithms for linear algebra [7, 8]. Approaches to mitigating this measurement problem include Pauli-term grouping [9–13] and improved allocation of measurements during parameterized circuit optimization [14]. While these techniques are helpful, the results of [6] suggest that they are insufficient for resolving the measurement problem. A more scalable approach is to use quantum algorithms which reduce the runtime of statistical estimation. Recent work [15] has developed a set of methods for *robust amplitude estimation*: the task of estimating amplitudes in the presence of device error. Such methods provide an interpolation between the performance of standard VQE and the performance of VQE in the fault-tolerant regime where quantum amplitude estimation is used. The concept of such an interpolation was introduced in the context of VQE [16] and more recently has been explored in the context of quantum algorithms for finance [17–19].

An outstanding question from the above theoretical work is: what are the limitations of robust amplitude

estimation when implemented with real quantum devices? In this paper we demonstrate that robust amplitude estimation already improves over standard VQE in the efficiency of estimating Pauli expectation values. Furthermore, we find that robust amplitude estimation yields a degree of error mitigation without requiring additional samples (as other error mitigation methods do). The most striking feature of this method is that it suggests the use of quite deep quantum circuits. While in many settings one may want to keep overall circuit error rates below 10% or even 1%, our work suggests that runtime is minimized when using circuits with an overall circuit error rate of around 40%. Through experiments using two qubits, we show that robust amplitude estimation yields a factor of four reduction in RMSE compared to standard sampling (as typically used in VQE), when spending the same allotted runtime. In addition to a reduction in RMSE, we also find an increase in accuracy through a ten-fold bias reduction, demonstrating the error-mitigating capabilities of this method.

Robust amplitude estimation uses enhanced sampling to speed up the estimation of expectation values on noisy quantum computers. A detailed description of enhanced sampling techniques can be found in the following reference [15]. In contrast to the standard sampling method used, for example, in VQE, enhanced sampling methods enable a reduction of estimation runtime proportional to improvements in the quality of the quantum hardware. We explain the simplest variant of these techniques, which is used in our experiments.

Many quantum algorithms require the estimation of quantities encoded as expectation values

$$\Pi = \langle A | P | A \rangle, \quad (1)$$

where $|A\rangle = A|0^n\rangle$ in which A is the ansatz circuit, and P is an n -qubit Hermitian operator with eigenvalues ± 1 . A straightforward estimation scheme entails repeatedly

* Current affiliation: IBM Yorktown Heights

† peter@zapatacomputing.com

preparing state $|A\rangle$ and measuring P (in the case of a Pauli operator, each qubit is measured in the corresponding Pauli basis and the parity of the outcomes gives the measured eigenvalue of P). The estimate is taken to be the sample mean of the ± 1 -valued outcomes. We refer to this strategy as *standard sampling* and note that this is typically used in VQE and other near-term quantum algorithms.

The essence of enhanced sampling is to gather data from quantum states whose expectation values are functions of the expectation value of interest. It modifies the standard sampling approach by using the following steps to generate measurement outcomes: prepare the ansatz state $|A\rangle = A|0^n\rangle$, apply L Grover layers $U = A(2|0\rangle\langle 0|^{\otimes n} - \mathbb{I})A^\dagger P$, and then measure the Pauli observable P . This sequence is depicted in Figure 1. Compared to standard sampling, the only substantially new operation introduced is the reflection about the initial state $R_0 = 2|0\rangle\langle 0|^{\otimes n} - \mathbb{I}$.

We estimate $\Pi = \langle A|P|A\rangle$ from a set of such measurement outcomes with varying L . The estimate can be inferred from this data using a maximum likelihood estimation process (see discussion around Eq. 4) based on the likelihood functions

$$\mathbb{P}(\pm 1|\Pi; L) = \frac{1}{2} (1 \pm T_{2L+1}(\Pi)), \quad (2)$$

where $T_m(x) = \cos(m \arccos(x))$ is the m -degree Chebyshev polynomial. Note that the expression for $\mathbb{P}(-1|\Pi; L)$ coincides with the Grover algorithm success probability due to the use of Grover layers U . Enhanced sampling reduces the runtime needed to obtain an accurate estimate by drawing measurement outcomes whose likelihoods depend sensitively on the parameter of interest.




FIG. 1: This figure depicts the operations for generating measurement outcomes used in enhanced sampling. A is the state preparation circuit, P is the observable of interest, and U is the Grover iterate comprised of $AR_0A^\dagger P$, with R_0 the reflection about the state $|0\rangle^{\otimes N}$.

In practice, quantum computation is subject to errors. These errors are caused by several sources, including decoherence to the thermal environment and limitations on the calibration of quantum gates. We can estimate expectation values despite such error by incorporating a model of the effect of error on the likelihood function into the inference process. If the model is sufficiently accurate and if we learn the model parameters to within a certain tolerance, then we may still be able to estimate the parameter of interest accurately. We adopt the exponential

decay model described in [15],

$$\mathbb{P}(\pm 1|\Pi, \lambda; L) = \frac{1}{2} \left(1 \pm e^{-(L+1/2)\lambda} T_{2L+1}(\Pi) \right), \quad (3)$$

where λ is a decay parameter roughly proportional to the error rate of the elementary gates on the quantum device and, again, $T_m(x) = \cos(m \arccos(x))$ is the m -degree Chebyshev polynomial. We emphasize that the model proposed here is an approximation to the actual relationship between the outcome likelihoods and the parameter of interest. Therefore, an important question addressed by this work is: how well does robust amplitude estimation perform in practice when using this approximate model of the likelihood function?

In each step of the inference process we update a prior distribution $p(\Pi, \lambda)$, according to the parity of a bitstring outcome $s = \pm 1$ generated from an L -layer enhanced sampling circuit and measurement, to obtain a posterior distribution

$$p(\Pi, \lambda|s) = \frac{\mathbb{P}(s|\Pi, \lambda; L)p(\Pi, \lambda)}{\int d\Pi d\lambda \mathbb{P}(s|\Pi, \lambda; L)p(\Pi, \lambda)}. \quad (4)$$

After many measurements and Bayesian updates we determine the values $\hat{\Pi}$ and $\hat{\lambda}$ which maximize the posterior distribution and use this $\hat{\Pi}$ as the estimate of the parameter of interest. The decay parameter λ is treated as a nuisance parameter [20], used as an intermediary for arriving at the maximum likelihood estimate (MLE) of Π . Enhanced sampling methods vary in the means by which L is chosen for each step and the way in which the Bayesian inference process is carried out numerically [15, 17, 21, 22].

To demonstrate value of the robust amplitude estimation method in a setting of interest, we selected the ansatz circuit of [23], which was used to calculate the ground state energy of the hydrogen molecule. The circuit is reproduced in Figure 2. For observables P , we restrict ourselves to Pauli strings X_0X_1, Y_0Y_1 , which appear in the two-qubit hydrogen molecule Hamiltonian [23].




FIG. 2: Circuit diagram of the VQE ansatz for molecular hydrogen we chose to use in our experiments. We chose the the circuit parameter to be $\theta = -6.0575$, which yields the ground state with respect to the molecular hydrogen Hamiltonian at 0.2 Angstroms. The expectation values of X_0X_1 and Y_0Y_1 are both -0.2238 for this state.

For each Pauli string, we generate parity measurement outcomes from the circuit of Figure 1 for a range of L , the number of Grover iterates. To generate the expectation value estimates, we process the measurement outcome data using Bayesian inference. We carry out inference from the sampled data using numerical updates

and maximum likelihood estimation (MLE) similar to [17, 21]. We implement MLE by keeping track of a discretization of the posterior distributions resulting from Eq. 4. After the data has been processed through these inference steps, we find a pair (λ, Π) that maximizes the posterior distribution. We describe this procedure in more detail. We define enhanced sampling circuits $\mathcal{C} = (\mathcal{C}_0, \mathcal{C}_1, \mathcal{C}_2, \dots, \mathcal{C}_{L_{max}})$ and sets of parity outcomes $\mathcal{M} = (\mathcal{M}_0, \mathcal{M}_1 \dots \mathcal{M}_{L_{max}})$, where \mathcal{C}_i is a circuit of the type in Figure 1 with i layers and \mathcal{M}_i are the corresponding set of ± 1 parity outcomes resulting from N_i total measurements and we set $N = \sum_{i=0}^{L_{max}} N_i$. The algorithm for processing this outcome data into estimates is as follows:

Maximum Likelihood Estimation Estimate the value of $\Pi = \langle A | P | A \rangle$ from a grid of candidate points

Input: $\mathcal{M} = (\mathcal{M}_0, \mathcal{M}_1, \dots, \mathcal{M}_{L_{max}})$ and prior distribution on a discrete grid $\mathcal{D} : (p, q) \rightarrow \mathbb{R}$


Output: (Π, λ)

- 1: Randomly sample, with duplicates, $M \leq N$ parity outcomes from \mathcal{M} to get (m_0, m_1, \dots, m_M) to create an ordered set \mathcal{R} and a corresponding ordered set of $\hat{\mathcal{T}} = (l_0, l_1, \dots, l_M)$ consisting of layer labels
 - 2: **for** $i = 0$ to $i = M$ **do**
 - 3: **if** $m_i = 1$ **then**
 - 3: $\mathcal{D}_{p,q} \leftarrow \mathcal{D}_{p,q} \times \frac{1}{2}(1 + e^{-(l_i+1/2)\lambda_p} T_{2l_i+1}(\Pi_q))$
 - 4: **else**
 - 4: $\mathcal{D}_{p,q} \leftarrow \mathcal{D}_{p,q} \times \frac{1}{2}(1 - e^{-(l_i+1/2)\lambda_p} T_{2l_i+1}(\Pi_q))$
 - 5: **end if**
 - 6: **end for**
 - 7: Find (Π_q^*, λ_p^*) such that $\mathcal{D}_{p,q}$ is maximized over all (p, q)
 - 8: $(\Pi, \lambda) \leftarrow (\Pi_q^*, \lambda_p^*)$
 - 9: **return** (Π, λ)
-


In our experiments we chose to use two different IBM devices: *ibmq-belem* and *ibmq-manila*. We varied the number of layers between 0 and 10. We generated a total of $N = 90112$ samples, with $N_L = 8192$ samples taken for each of the 11 layer numbers L .

We post-processed this data according to the maximum likelihood estimation algorithm shown above with $M = 1000$ samples per estimate. We run the maximum likelihood estimation algorithm 32 times, each time choosing a random sample of outcome data. We choose 32 trials because the number of independent samples is 32 times the number of samples used per trial, making each trial nearly independent.

Figure 3 shows the estimated expectation value of $\langle Y_0 Y_1 \rangle$ and $\langle X_0 X_1 \rangle$ as we increase L_{max} , the maximum number of Grover iterates used. We observe a reduction in the RMSE as we increase the layers up to $L_{max} = 3$ for *ibmq-belem*, beyond which there is no substantial reduction in RMSE, while for *ibmq-manila* the improvement is observed until $L_{max} \approx 5$. We expect that in the regime of $L_{max} > 3$ for *ibmq-belem* and $L_{max} > 5$ for *ibmq-manila*, the deep circuits used have accrued too much error to yield valuable information.



(a) The two plots compare the estimates of $\langle Y_0 Y_1 \rangle$ and $\langle X_0 X_1 \rangle$, for different values of L_{max} . In the plots, each point is the mean among the 1000 estimates while the error bars indicate the standard deviation among these estimates. The dashed line represents the exact value of $\langle Y_0 Y_1 \rangle$ and $\langle X_0 X_1 \rangle$.



(b) In these plots we study the change in the RMSE of the estimates while varying L_{max} . For *ibmq-belem* devices, we see that the RMSE improves substantially until about $L_{max} \approx 3$ after which we see now substantial gain in sampling power on the other hand *ibmq-manila* sees gains until $L_{max} \approx 5$. The error bars are computed as the square root of the standard deviation of the squared error from each trial, divided by the square root of the number of trials (i.e. 32).

FIG. 3: Performance of robust amplitude estimation for estimating Π for various L_{max} , the maximum number of Grover iterates used per circuit. Note that we do not include $L_{max} = 0$ because we are unable to resolve two unknown parameters (i.e. Π and λ) with a single parameter frequency estimate from $L = 0$ samples.

In Figure 4 we demonstrate the surprising error mitigating capability of robust amplitude estimation. For each value of L_{max} , we plot the corresponding estimated bias and standard deviation of the expectation value estimates. The substantial error mitigation is indicated by a nearly ten-fold reduction in the bias as the maximum layer number is increased from $L_{max} = 1$. We expect that the use of varying layer number in inferring the expectation value helps to average out the contributions from coherent error that differs from layer to layer.




FIG. 4: We show the change in bias and standard deviation, σ , in the $\langle Y_0 Y_1 \rangle$ and $\langle X_0 X_1 \rangle$ expectation value estimates as the maximum number of Grover iterates is increased from 1 to 10. The RMSE is the distance from the origin. This plot shows that, in addition to an improvement in precision, RAE yields a reduction in bias indicating a degree of error mitigation.

These results lead to two main observations:

- (i) Robust amplitude estimation can yield estimates with lower RMSE than standard sampling.
- (ii) Robust amplitude estimation exhibits a hitherto unforeseen error mitigation capability.

It is worth remarking that these two improvements are achieved in parallel; this contrasts with traditional error mitigation methods (e.g. zero-noise extrapolation [24] and readout error mitigation schemes like [25]), where reduction in error comes at the cost of increased samples and run-time without the reduction in precision.

The surprising conclusion is the following: to reduce runtimes and error in estimation subroutines (e.g. as used in VQE) with near-term quantum devices we should employ deeper and noisier quantum circuits. These deeper quantum circuits use quantum amplification to increase the statistical power of the measurement outcome data. This counter-intuitive conclusion can be understood in the following manner: a crucial ingredient in the noisy implementation of our algorithm is the likelihood function (c.f. Eq. 3) which incorporates a model of how noise (captured by the nuisance parameter) impacts the likelihoods of outcome data. The key point is that as long as this model is not far from the true likelihood function, the inference procedure can still yield accurate estimates, where the noise merely slows the rate of information gain as opposed to biasing the estimates. This we believe highlights the importance of noise-tailoring for the NISQ era.

We have shown that robust amplitude estimation may increase precision and accuracy. We now investigate whether robust amplitude estimation achieves a better

RMSE than standard VQE estimation when both methods are given the same allotted runtime. We will restrict this exploration to the more-performant of the two IBM devices, *ibmq_manila*. Furthermore, we will only consider the case of estimating $\langle X_0 X_1 \rangle$ and not $\langle Y_0 Y_1 \rangle$ because, according to Fig. 3, RAE performed better for this case. We expect that noise tailoring methods could help to remove the bias in the RAE estimates of $\langle Y_0 Y_1 \rangle$, but we leave this to future investigation. In the standard sampling setting, we generate estimates assuming that $\lambda = 0$, corresponding to the standard approach used in VQE. To fix the runtimes to be equal, we account for the fact that the circuit depths (and therefore circuit times) are different in the two settings. While each circuit in the VQE setting uses just one query to the ansatz circuit, each enhanced sampling circuit makes $(2L + 1)$ queries to the ansatz. Furthermore, it makes L uses of the phase oracle, which we take to be roughly half the depth of the ansatz (in this two-qubit setting). To increase the rate of information gain in RAE, we choose to only draw samples from circuits whose likelihood functions yield large Fisher information per time:

$$\frac{\mathcal{I}}{T} = \frac{1}{x} \frac{e^{-\lambda x} \sin^2(x\theta)}{1 - e^{-\lambda x} \cos^2(x\theta)}, \quad (5)$$

where $x = 2L + 1$ and $\theta = \arccos \Pi$. Setting $\lambda = 0.08$ and $\cos \theta = -0.22$, which are coarse approximations to the values resulting from maximum likelihood estimation, the Fisher information per time peaks around $L = 1$ and $L = 7$ as shown in Figure 5. We choose to infer Π using samples from $L = 1, 5, 6, 7$; We choose the total number of shots such that the total circuit time is equal to that of 12875 shots with $L = 0$ (i.e. the VQE setting). By taking $L = 1, 5, 6$, and 7 , the corresponding depths in units of the ansatz depth are 3.5, 13.5, 16, and 18.5. Thus, the average circuit depth (and, therefore, circuit time) will be 12.875 times that of the ansatz. Thus, to achieve the same runtime between the settings of VQE and robust amplitude estimation ($L = 1, 5, 6, 7$), we allow $M = 12875$ and $M = 1000$ samples, respectively.

The results are reported in Table I; we find that robust amplitude estimation yields roughly a four-fold reduction in RMSE and two-fold improvement in precision while costing roughly the same runtime. Furthermore, there is roughly a ten-fold improvement in the bias, indicating a substantial degree of error mitigation achieved by robust amplitude estimation. We note that this run-time costing has ignored the circuit latency times, which, if accounted for, would lead to an even further separation in performance between the two approaches.

In this work we have explored an approach to solving the so-called “measurement problem”, which plagues many near-term quantum algorithms, including VQE. It is believed that achieving quantum advantage with near-term quantum computers will require solving this problem [6]. The measurement problem can be attributed to




FIG. 5: This plot shows the Fisher information per unit time for different numbers of Grover iterates. We use the true value of $\Pi = -0.22$ and $\lambda = 0.08$, which are rough approximations to the values obtained through Bayesian inference. The maximum Fisher information rate is achieved for $L = 7$, which is close to satisfying the rule of thumb for setting the optimal layer number: $2L + 1 \rightarrow \frac{1}{\lambda}$ (c.f. [15]). We note that the plot provides an example of a “dead spot” as introduced in [15]: the exponential decay factor, introduced to model error in the quantum circuit, causes certain Grover circuits to yield unfavorable Fisher information.

Method	RMSE	σ (Prec.)	Bias (Acc.)	Samples
VQE	0.025(2)	0.011(1)	0.022(2)	12,875
RAE	0.0045(6)	0.0043(5)	0.0012(8)	1,000

TABLE I: This table compares the performance of the variational quantum eigensolver (VQE) and robust amplitude estimation (RAE) when given the same allotted runtime. We observe a factor of four improvement in RMSE, a factor of two improvement in precision, and a factor of ten improvement in accuracy (as measured by the bias) when estimating $X_0 X_1$ on *ibmq_manila*. Here, each RMSE, σ , and bias are estimated using 19 independent trials for VQE and 32 independent trials for RAE.

the low statistical power of typical prepare-and-measure (referred to as “standard sampling”) estimation methods. Fortunately, quantum computing enables methods to increase the statistical power of estimation tasks through *quantum amplification*, providing a path to solving the

measurement problem.

We have investigated how recently-introduced techniques [15] involving quantum amplification, which we refer to as *robust amplitude estimation*, stand to improve the performance of estimating expectation values with imperfect quantum computers. The primary open question from previous work was: how do such methods perform in practice? This work may be seen as a step towards assessing robust amplitude estimation *in practice* as a viable solution to the measurement problem.

Using several different publicly-available IBM devices, we implemented both robust amplitude estimation and estimation with standard sampling for estimating two-qubit Pauli expectation values. We found that robust amplitude estimation gives a substantial improvement in two distinct ways: reduction in runtime and error mitigation.

We hope that our findings encourage practitioners of near-term quantum algorithms to begin adopting enhanced sampling and robust amplitude estimation as valuable tools for their quantum computations. As the demand for using quantum devices increases, quantum compute time will become more costly. The runtime savings afforded by robust amplitude estimation translate directly to a monetary cost savings.

This work leaves a number of directions open for future investigation. It remains to develop a rigorous analysis and model for the performance of these techniques as the system size is scaled to larger qubit counts. Improving over the algorithm performance model in [15], such a model would be very valuable for carrying out resource estimations and assessing prospects for quantum advantage. Another important direction is to develop methods for estimating multiple parameters in parallel, which is a feature of standard sampling as used in VQE. In future work we plan to incorporate noise-tailoring techniques, such as randomized compiling [26], to make the actual likelihood function closer to the likelihood function model, improving the performance of estimation. With each of these directions in mind and with the promising findings on robust amplitude estimation, we hope this work will encourage researchers to develop methods that are necessary for realizing useful quantum computation.

Acknowledgements: The authors thank Jérôme Gonthier, Artur Izmaylov, Jens Eisert, Yudong Cao, Michał Stechły, and Dax Koh for helpful feedback on the manuscript. Experiments were run using the Zapata Computing Platform Orquestra™.

-
- [1] Peter JJ O’Malley, Ryan Babbush, Ian D Kivlichan, Jonathan Romero, Jarrod R McClean, Rami Barends, Julian Kelly, Pedram Roushan, Andrew Tranter, Nan Ding, et al. Scalable quantum simulation of molecular energies. *Physical Review X*, 6(3):031007, 2016.
- [2] Frank Arute, Kunal Arya, Ryan Babbush, Dave Bacon, Joseph C Bardin, Rami Barends, Rupak Biswas, Sergio

- Boixo, Fernando GSL Brandao, David A Buell, et al. Quantum supremacy using a programmable superconducting processor. *Nature*, 574(7779):505–510, 2019.
- [3] Yunseong Nam, Jwo-Sy Chen, Neal C Pienti, Kenneth Wright, Conor Delaney, Dmitri Maslov, Kenneth R Brown, Stewart Allen, Jason M Amini, Joel Apisdorf, et al. Ground-state energy estimation of the water

- molecule on a trapped-ion quantum computer. *npj Quantum Information*, 6(1):1–6, 2020.
- [4] Manuel S Rudolph, Ntwali Toussaint Bashige, Amara Katarbarwa, Sonika Johr, Borja Peropadre, and Alejandro Perdomo-Ortiz. Generation of high resolution handwritten digits with an ion-trap quantum computer. *arXiv preprint arXiv:2012.03924*, 2020.
- [5] Alberto Peruzzo, Jarrod McClean, Peter Shadbolt, Man-Hong Yung, Xiao-Qi Zhou, Peter J Love, Alán Aspuru-Guzik, and Jeremy L O’Brien. A variational eigenvalue solver on a photonic quantum processor. *Nature communications*, 5(1):1–7, 2014.
- [6] Jérôme F Gonthier, Maxwell D Radin, Corneliu Buda, Eric J Daskocil, Clena M Abuan, and Jhonathan Romero. Identifying challenges towards practical quantum advantage through resource estimation: the measurement roadblock in the variational quantum eigensolver. *arXiv preprint arXiv:2012.04001*, 2020.
- [7] Marcello Benedetti, Delfina Garcia-Pintos, Oscar Perdomo, Vicente Leyton-Ortega, Yunseong Nam, and Alejandro Perdomo-Ortiz. A generative modeling approach for benchmarking and training shallow quantum circuits. *npj Quantum Information*, 5(1):1–9, 2019.
- [8] Rolando D Somma and Yiğit Subaşı. Complexity of quantum state verification in the quantum linear systems problem. *PRX Quantum*, 2(1):010315, 2021.
- [9] Artur F. Izmaylov, Tzu-Ching Yen, Robert A. Lang, and Vladyslav Verteletskyi. Unitary partitioning approach to the measurement problem in the variational quantum eigensolver method. *Journal of Chemical Theory and Computation*, 16(1):190–195, 2020. PMID: 31747266.
- [10] Andrew Zhao, Andrew Tranter, William M Kirby, Shu Fay Ung, Akimasa Miyake, and Peter J Love. Measurement reduction in variational quantum algorithms. *Physical Review A*, 101(6):062322, 2020.
- [11] William J Huggins, Jarrod R McClean, Nicholas C Rubin, Zhang Jiang, Nathan Wiebe, K Birgitta Whaley, and Ryan Babbush. Efficient and noise resilient measurements for quantum chemistry on near-term quantum computers. *npj Quantum Information*, 7(1):1–9, 2021.
- [12] Ophelia Crawford, Barnaby van Straaten, Daochen Wang, Thomas Parks, Earl Campbell, and Stephen Brierley. Efficient quantum measurement of pauli operators in the presence of finite sampling error. *Quantum*, 5:385, 2021.
- [13] Alexis Ralli, Peter J Love, Andrew Tranter, and Peter V Coveney. Implementation of measurement reduction for the variational quantum eigensolver. *Physical Review Research*, 3(3):033195, 2021.
- [14] Ryan Sweke, Frederik Wilde, Johannes Meyer, Maria Schuld, Paul K Fährmann, Barthélémy Meynard-Piganeau, and Jens Eisert. Stochastic gradient descent for hybrid quantum-classical optimization. *Quantum*, 4:314, 2020.
- [15] Guoming Wang, Dax Enshan Koh, Peter D Johnson, and Yudong Cao. Minimizing estimation runtime on noisy quantum computers. *PRX Quantum*, 2(1):010346, 2021.
- [16] Daochen Wang, Oscar Higgott, and Stephen Brierley. Accelerated variational quantum eigensolver. *Physical review letters*, 122(14):140504, 2019.
- [17] Rudy Raymond Tomoki Tanaka Tamiya Onodera & Naoki Yamamoto Yohichi Suzuki, Shumpei Uno. Amplitude estimation without phase estimation. *Quantum Information Processing volume*, 19(75), jan 2020.
- [18] Javier Alcazar, Andrea Cadarso, Amara Katarbarwa, Marta Mauri, Borja Peropadre, Guoming Wang, and Yudong Cao. Quantum algorithm for credit valuation adjustments. *arXiv preprint arXiv:2105.12087*, 2021.
- [19] Tudor Giurgica-Tiron, Sonika Johri, Iordanis Kerenidis, Jason Nguyen, Neal Pisenti, Anupam Prakash, Ksenia Sosnova, Ken Wright, and William Zeng. Low depth amplitude estimation on a trapped ion quantum computer. *arXiv preprint arXiv:2109.09685*, 2021.
- [20] Richard Royall. On the probability of observing misleading statistical evidence. *Journal of the american statistical association*, 95(451):760–768, 2000.
- [21] W.K. Tham Eric G. Brown, Oktay Goktas. Quantum amplitude estimation in the presence of noise. <https://arxiv.org/abs/2006.14145>, jun 2020.
- [22] Tudor Giurgica-Tiron, Iordanis Kerenidis, Farrokh Labib, Anupam Prakash, and William Zeng. Low depth algorithms for quantum amplitude estimation. *arXiv preprint arXiv:2012.03348*, 2020.
- [23] P. J. J. O’Malley, R. Babbush, I. D. Kivlichan, J. Romero, J. R. McClean, R. Barends, J. Kelly, P. Roushan, A. Tranter, N. Ding, B. Campbell, Y. Chen, Z. Chen, B. Chiaro, A. Dunsworth, A. G. Fowler, E. Jeffrey, E. Lucero, A. Megrant, J. Y. Mutus, M. Neeley, C. Neill, C. Quintana, D. Sank, A. Vainsencher, J. Wenner, T. C. White, P. V. Coveney, P. J. Love, H. Neven, A. Aspuru-Guzik, and J. M. Martinis. Scalable quantum simulation of molecular energies. *Phys. Rev. X*, 6:031007, Jul 2016.
- [24] Tudor Giurgica-Tiron, Yousef Hindy, Ryan LaRose, Andrea Mari, and William J Zeng. Digital zero noise extrapolation for quantum error mitigation. In *2020 IEEE International Conference on Quantum Computing and Engineering (QCE)*, pages 306–316. IEEE, 2020.
- [25] George S Barron and Christopher J Wood. Measurement error mitigation for variational quantum algorithms. *arXiv preprint arXiv:2010.08520*, 2020.
- [26] Joel J Wallman and Joseph Emerson. Noise tailoring for scalable quantum computation via randomized compiling. *Physical Review A*, 94(5):052325, 2016.