

CAFQA: A classical simulation bootstrap for variational quantum algorithms

Gokul Subramanian Ravi*
University of Chicago

Pranav Gokhale
Super.tech (a division of ColdQuanta)

Yi Ding
MIT

William M. Kirby
Tufts University

Kaitlin N. Smith
University of Chicago

Jonathan M. Baker
University of Chicago

Peter J. Love
Tufts University

Henry Hoffmann
University of Chicago

Kenneth R. Brown
Duke University

Frederic T. Chong
University of Chicago
Super.tech (a division of ColdQuanta)

ABSTRACT

Classical computing plays a critical role in the advancement of quantum frontiers in the NISQ era. In this spirit, this work uses classical simulation to bootstrap Variational Quantum Algorithms (VQAs). VQAs rely upon the iterative optimization of a parameterized unitary circuit (ansatz) with respect to an objective function. Since quantum machines are noisy and expensive resources, it is imperative to classically choose the VQA ansatz initial parameters to be as close to optimal as possible to improve VQA accuracy and accelerate their convergence on today's devices.

This work tackles the problem of finding a good ansatz initialization, by proposing CAFQA, a Clifford Ansatz For Quantum Accuracy. The CAFQA ansatz is a hardware-efficient circuit built with only Clifford gates. In this ansatz, the parameters for the tunable gates are chosen by searching efficiently through the Clifford parameter space via classical simulation. The resulting initial states always equal or outperform traditional classical initialization (e.g., Hartree-Fock), and enable high-accuracy VQA estimations. CAFQA is well-suited to classical computation because: a) Clifford-only quantum circuits can be exactly simulated classically in polynomial time, and b) the discrete Clifford space is searched efficiently via Bayesian Optimization.

For the Variational Quantum Eigensolver (VQE) task of molecular ground state energy estimation (up to 18 qubits), CAFQA's Clifford Ansatz achieves a mean accuracy of nearly 99% and recovers as much as 99.99% of the molecular correlation energy that is lost in Hartree-Fock initialization. CAFQA achieves mean accuracy

improvements of 6.4x and 56.8x, over the state-of-the-art, on different metrics. The scalability of the approach allows for preliminary ground state energy estimation of the challenging chromium dimer (Cr_2) molecule. With CAFQA's high-accuracy initialization, the convergence of VQAs is shown to accelerate by 2.5x, even for small molecules.

Furthermore, preliminary exploration of allowing a limited number of non-Clifford (T) gates in the CAFQA framework, shows that as much as 99.9% of the correlation energy can be recovered at bond lengths for which Clifford-only CAFQA accuracy is relatively limited, while remaining classically simulable.

CCS CONCEPTS

• Computer systems organization → Quantum computing.

KEYWORDS

quantum computing, variational quantum algorithms, chemistry, clifford, bayesian optimization, noisy intermediate-scale quantum, variational quantum eigensolver

ACM Reference Format:

Gokul Subramanian Ravi, Pranav Gokhale, Yi Ding, William M. Kirby, Kaitlin N. Smith, Jonathan M. Baker, Peter J. Love, Henry Hoffmann, Kenneth R. Brown, and Frederic T. Chong. 2022. CAFQA: A classical simulation bootstrap for variational quantum algorithms. In *Proceedings of 28th Conference on Architectural Support for Programming Languages and Operating Systems (ASPLOS)*. ACM, New York, NY, USA, 15 pages. <https://doi.org/10.1145/nnnnnnnn.nnnnnnnn>

1 INTRODUCTION

Quantum Computing in the NISQ era: Quantum computing (QC) is a revolutionary computational model to solve certain classically intractable problems and is projected to give QCs a significant advantage in cryptography [65], chemistry [36], optimization [47] and machine learning [8]. In the ongoing Noisy Intermediate-Scale Quantum (NISQ) era, we expect to work with quantum machines which comprise hundreds to thousands of imperfect qubits [57]. On the one hand, NISQ era machines will be unable to execute large-scale quantum algorithms like Shor Factoring [65] and Grover Search [30], which would require error correction comprised of

*Correspondence: gravi@uchicago.edu

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

ASPLOS, 2023, Vancouver, Canada

© 2022 Association for Computing Machinery.

ACM ISBN 978-x-xxxx-xxxx-x/YY/MM...\$15.00

<https://doi.org/10.1145/nnnnnnnn.nnnnnnnn>




Figure 1: Advancing the NISQ frontiers with error mitigation, and classical support (eg., CAFQA).

millions of qubits to create fault-tolerant quantum systems [52]. On the other hand, a variety of error mitigation techniques [7, 10, 15, 17, 24, 41, 61, 66, 69, 72, 78] have been proposed that have improved execution fidelity on today’s quantum devices. However, the resulting fidelity is still insufficient for most real-world use cases.

Advancing NISQ with classical support: There has been a recent impetus toward classical computing support to boost NISQ applications / devices to the realm of real-world applicability. These include compiler level optimizations [48, 49, 58, 71], improved classical optimizers [39], circuit cutting with classical compensation [20, 70, 80] etc. We are still in the early days of exploring this synergistic quantum-classical paradigm. There is tremendous potential for sophisticated application-specific classical bootstrapping to advance the NISQ frontiers, and CAFQA is one such approach. An illustration of advancing the NISQ frontiers towards real-world applicability is shown in Fig.1.

Variational Quantum Algorithms: Variational quantum algorithms (VQAs) are expected to be a good match for NISQ machines. This class of algorithms has a wide range of applications, such as the estimation of electronic energy of molecules [55], MAXCUT approximation [47]. The quantum circuit for a VQA is parameterized by a list of angles which are optimized by a classical optimizer over many iterations towards a specific target objective which is representative of the VQA problem. VQAs are more suitable for today’s quantum devices because these algorithms adapt to the characteristics and noise profile of the quantum machine on [44, 55]. Unfortunately, VQA accuracy obtained on today’s NISQ machines, even with error mitigation, is often considerably far from the stringent accuracy requirements in fields such as molecular chemistry, especially as we scale to larger problem sizes [36, 58, 78].

Aiding VQAs in the NISQ Era: For NISQ VQAs to progress towards real-world applicability, it is imperative to classically choose a VQA’s parameterized circuit (ansatz) wisely and its initial parameters to be as close to optimal as possible, prior to quantum exploration. This would improve accuracy and accelerate convergence of the algorithm on the noisy quantum device [43, 79]. Suitable ansatz circuits for today’s devices, referred to as “hardware efficient ansatz” [36], are often application-agnostic and can especially




Figure 2: Figure illustrates the VQA(E) state space explored by different techniques. Hartree-Fock (HF), only explores the yellow computational basis (i.e., classical bits) space. CAFQA goes further by exploring the Clifford space (blue) for higher accuracy. Both of these are classically simulable. Classical exploration can be extended to incorporate a few T gates (orange). The rest of the state space is shown in green - this space cannot be efficiently simulated classically and requires quantum exploration. By searching efficiently through the classically simulable Clifford space, CAFQA provides a good VQA ansatz initialization for quantum exploration.

benefit from a wise choice of initial parameters, but these can be difficult to estimate classically.

Classical simulation support for VQAs with Cliffords: This work helps to initialize the VQA ansatz using classical simulation. In general, classical simulation of quantum tasks is not a scalable solution, primarily only suited to trivial quantum problems, and is, in fact, the motivation for quantum machines. An exception to the above is the classical simulation of the Clifford space. Circuits made up of only Clifford operations can be exactly simulated in polynomial time [27]. Clifford operations do not provide a universal set of quantum gates - hence, the stabilizer states produced by Clifford-only circuits are limited in how effectively they can explore the quantum space of a given problem such as those targeted by VQAs. However, exploring the Clifford space of the VQA problem through ideal classical simulation can potentially find good noise-free initial states, which is particularly beneficial in the NISQ era. Note: in addition to Clifford gates, it is possible for a small number of T gates to also be efficiently classically simulated [12]. While we primarily focus on the Clifford space in this work, we show preliminary results for beyond-Clifford exploration in Section 8.

CAFQA: This work tackles the challenge of finding initial ansatz parameters by proposing CAFQA, a Clifford Ansatz For Quantum Accuracy. The CAFQA Clifford Ansatz is a hardware-efficient parameterizable circuit that is parameterized with only Clifford gates.

In this ansatz, the initial parameters for the tunable gates are chosen by searching efficiently through the Clifford parameter space via classical simulation, thereby producing a suitable stabilizer state. The proposed approach is attractive for multiple reasons: ① the Clifford-initialized ansatz produces stabilizer states that perform equal to or better than the traditional classical approach of finding a suitable computational basis state (e.g., Hartree-Fock [32]) because (i) it can explore a larger state space and (ii) the stabilizer can have direct chemical relevance in some molecular systems; ② the Clifford-only quantum circuits can be perfectly simulated in polynomial (quadratic or even linear) time on classical computers [27]; ③ the produced initial states are obtained ideally, since classical simulation is noise-free; ④ the discrete Clifford space is searched efficiently via Bayesian Optimization using a random forest surrogate model and a greedy acquisition function; ⑤ while the Clifford space is significantly smaller compared to the entire quantum space, the stabilizer states produced from the optimal Clifford parameters are able to achieve solutions of high accuracy even prior to execution/exploration on a quantum device; and finally ⑥ the selected ansatz can then be tuned over the entire parameter space on a quantum device, allowing for accelerated accurate convergence on NISQ devices (and beyond). Fig. 2 provides a break down of the VQA parameter space and CAFQA’s scope.

Key CAFQA results:

① For the VQE task of ground state energy estimations of molecular systems up to 18 qubits, CAFQA’s Clifford Ansatz is observed to achieve a mean accuracy of near 99% and is able to recover up to 99.99% of the molecular correlation energy lost in state-of-the-art Hartree-Fock initialization. CAFQA achieves mean accuracy improvements over the state-of-the-art of 6.4x when averaged over all bond lengths and 56.8x at highest bond lengths (maximum of $3.4 \times 10^5 \times$).

② Quantum exploration post CAFQA initialization can lead to faster and highly accurate VQA convergence, even on reasonably noisy quantum machines — we show 2.5x faster convergence compared to HF for a small molecule. Greater benefits can be expected for larger problem sizes, which can be usefully evaluated when NISQ machines improve.

③ The scalability of the approach allows for accurate ansatz initialization for ground state energy estimation of the challenging Chromium dimer (often considered a benchmark for variational quantum advantage) with greater than Hartree-Fock accuracy.

④ Preliminary exploration of allowing a very limited number of non-Clifford (T) gates in the CAFQA framework shows that as much as 99.9% of the correlation energy can be recovered at bond lengths for which Clifford-only CAFQA accuracy is relatively limited, while remaining classically simulable

Key CAFQA insights:

① CAFQA uses classical simulation to explore the Clifford space of a VQA problem and produces high accuracy VQA ansatz initialization, considerably outperforming the state-of-the-art.

② CAFQA’s benefits are especially significant because it is classically simulable, it searches the search space efficiently, and its evaluations are ideal.

③ CAFQA highlights the potential for quantum inspired classical techniques as well as a synergistic quantum-classical paradigm, to boost NISQ-era quantum computing (with focus on VQA) towards real world applicability.

2 BACKGROUND AND MOTIVATION

2.1 VQAs in the NISQ era

VQE: While CAFQA is suited widely across variational algorithms (eg., QAOA [22]), in this paper we primarily focus on the Variational Quantum Eigensolver (VQE) [55]. VQE is used to estimate an upper bound on the ground state energy of a Hamiltonian. Here, a Hamiltonian is a mathematical representation of some problem from, say, optimization or molecular chemistry, and is a linear combination of multiple Pauli terms. For example, a 4-qubit Hamiltonian could be $[H = 0.1 * XYXY + 0.5 * IZZI]$. VQE tries to find suitable parameters for an appropriately chosen parameterized circuit (ansatz) such that the expectation value of the target Hamiltonian is minimized. At a high level, VQE can be conceptualized as a repetitive “classical guess” + “quantum check” algorithm [25]. The check stage involves the preparation of a quantum state corresponding to the guess. This preparation stage is done in polynomial time on a quantum computer, but would incur exponential cost in general on a classical computer. This contrast gives rise to a potential quantum speedup for VQE [26]. In chemistry, VQE is a critical step in computing the energy properties of molecules and materials. While conventional computational chemistry provides methods to approximate such properties, they can lack sufficient accuracy in molecular systems due to an inadequate treatment of the correlations between constituent electrons. These interactions require computation that scales exponentially in the size of the system [73, 81].

NISQ era accuracy: Estimating the VQE global optimum with high accuracy has proven challenging in the NISQ era even with sophisticated optimizers, a well-chosen ansatz, and error mitigation [7, 10, 15, 58, 61, 69, 73, 78]. As an example, ground state energy estimation of molecules (the energy required to break a molecule into its sub-atomic components), a key use case for VQE, requires energy estimates with an estimation error of less than 1.6×10^{-3} Hartree, or what is known as “chemical accuracy” [56], for applicability in understanding chemical reactions and their rates. Unfortunately, for instance, previous work on the estimation of ground state energy of BeH₂ on a superconducting transmon machine resulted in an error greater than 10^{-1} Hartree, which is roughly 100x worse than the required accuracy [36]. Considering the significant disparity between NISQ VQA accuracy and real world requirements, it is imperative to aid VQA to the best extent possible.

2.2 VQA Ansatz and its Initialization

Ansatz: An ansatz is a parameterized circuit which is used to explore the quantum Hilbert space of the target VQA Hamiltonian, to find its ground state energy. An ansatz with parameterized gate rotation angles is shown in Fig. 3. Many ansatz structures are suitable for VQAs. In the context of VQE for molecular chemistry, the Unitary Coupled Cluster Single-Double (UCCSD) ansatz is considered the gold standard [26, 60]. Unfortunately, the UCCSD ansatz is generally of considerable circuit depth, making it less suitable for today’s NISQ machines, except for very small molecules such as H₂.




Figure 3: A Clifford Ansatz, a circuit with only Clifford gates. In this example, all fixed components are CXs and the tunable rotational gate angles are multiples of $\pi/2$.

More suitable to the NISQ-era are hardware-efficient ansatz like the SU2 [1] which are low depth parameterized circuits (but have Hilbert space coverage limitations [34, 73]). Fundamentally, this ansatz is constructed by repeating blocks of parameterized single-qubit rotation gates and ladders of entangling gates [73] (Fig.3). CAFQA builds atop a traditional hardware efficient ansatz [36]. Suitability to other ansatz structures is discussed in Section 8.

Optimization surface: While a good choice of classical optimizer improves VQA convergence [39], VQAs can: a) have complex optimization surfaces and b) suffer from a barren plateau problem. The optimization contour worsens as the noise and complexity of the problem increases in relation to the increase in the depth of the circuit, the number of parameters and the spread of the entanglement [39]. The barren plateau is the phenomenon in which the gradients of the VQE parameters vanish exponentially. While barren plateaus can become a critical issue for a variety of reasons [13, 42, 43, 54, 75], in the context of this work, they can become significant in the presence of noise [79] and with poor (random) ansatz initialization [43]. Thus, well-chosen initialization of the ansatz can help avoid barren plateaus and effects of noise, and therefore enable fast accurate convergence on the VQA problem.

Hartree-Fock initialization: A popular and simple approach to construct a fair initial state for quantum systems is derived from Hartree-Fock (HF) theory [32]. Although such approximation / optimization problems are generically hard, HF usually rapidly converges to good solutions, especially for closed-shell molecules at equilibrium geometries [2]. HF yields an initial state that has no entanglement between the electrons (i.e., simply a bitstring of 0s and 1s on the circuit’s qubits). HF assumes that each electron’s motion can be described as a stand-alone particle function, independent of the instantaneous motion of other electrons. In doing so, HF neglects the correlation between electrons, which is where classical computing is limited in solving such molecular chemistry problems. Therefore, although HF has reasonable accuracy for many molecules, it is generally insufficient to make highly accurate quantitative predictions [3]. Thus, its usefulness as a suitable initialization on today’s very noisy quantum devices is limited - there is

too much ground left for the quantum device to cover, which is challenging considering the complex noisy optimization surface and barren plateaus as described earlier. Therefore, initialization with greater accuracy, especially for strongly correlated systems and/or away from equilibrium geometry, necessitates quantum states that go beyond HF.

2.3 Clifford Circuits

Classical simulation of quantum problems usually requires exponential resources (otherwise, the need for quantum computers is obviated). Even with high-performance supercomputers, simulation is restricted to under 100 qubits [9, 16, 35, 46, 73].

However, not all simulations are non-scalable. The Gottesman-Knill theorem states that “Any quantum computer performing only: a) Clifford group gates, b) measurements of Pauli group operators, and c) Clifford group operations conditioned on classical bits, which may be the results of earlier measurements, can be perfectly simulated in polynomial time on a probabilistic classical computer” [27].

While the Clifford group operations and Pauli group measurements do not provide a universal set of quantum gates, there are quantum domains that have applications focused on the Clifford-space including quantum networks [77], error correction codes [59], teleportation [28] and error mitigation [15, 67].

CAFQA explores the benefits of Clifford-only circuits as an ansatz for variational algorithms. An example is shown in Fig.3. Extending beyond Cliffords is discussed in Section 8.

3 CAFQA PROPOSAL

Fig.4 provides an illustrative overview of how CAFQA complements traditional VQA tuning. CAFQA is illustrated in the red box and is discussed below.

① CAFQA begins with a parameterized circuit in which all fixed gates are Clifford. This is usually the case with hardware-efficient ansatz, as described in Section 2.2. Focusing on a hardware-efficient ansatz is justified considering that other ansatz options are generally less suitable to noisy execution on today’s NISQ devices. However, extensions are discussed in Section 8.

② Given this parameterized circuit, CAFQA performs a discrete search over the tunable circuit parameters. The tunable search space is limited to angles which make the tunable gates Clifford. Extensions discussed in Section 8.

③ Since both the fixed gates as well as the tunable gates are Clifford, the resulting circuit in each iteration of the tuning process can be simulated classically, even as the size of the circuits grow (as discussed in Section 2.3).

④ Simulating the ansatz circuits corresponding to the Hamiltonian and measuring the expectation produces the objective function value for the iterative tuning process.

⑤ For molecular chemistry, if any electron and spin preservation constraints have to be imposed on the problem they can be added to the Hamiltonian [63] or directly to the objective function - CAFQA uses the latter. More in Section 7.1.

⑥ Since simulations are performed classically, they are free of noise, thus having the potential to eliminate a considerable portion of the noise impact that variational tuning on the real quantum device could suffer, i.e., noise-induced barren plateaus [79] etc.




Figure 4: The red box on the left shows the CAFQA framework. CAFQA performs ansatz parameter tuning inspired by traditional VQA. But CAFQA’s tuning is suitable for classical compute since it restricts the search space to the Clifford space alone (in an ansatz that has its fixed components to also be Clifford, often the case with a hardware-efficient ansatz). The search space is discrete and is searched with Bayesian Optimization. Although the search space is limited, its evaluation is classically efficient and noise-free. On the other hand, traditional quantum variational tuning, while scalable and suited to the entire quantum space, is extremely noisy in the NISQ era. Once CAFQA finds a suitable Clifford initialization, traditional VQA tuning is performed (blue box / right), leading to faster and more accurate convergence.

⑦ Also worth noting is that only one-shot simulation is required for each Pauli term, since the expectation value produced by each term is strictly +1, -1, or 0 for stabilizer states (i.e., for Clifford circuits) [51].

⑧ The search is continued until the convergence of the minimum value obtained or for a specific tuning budget.

⑨ The resulting circuit, with Clifford parameters corresponding to the minimum objective function value observed, is the Clifford ansatz and is then ready for traditional VQA optimization.

⑩ Subsequent VQA tuning on the quantum device is noisy but is able to explore the entire quantum space allowed by the ansatz. The initial state produced by CAFQA can enable faster and more accurate convergence of traditional VQA.

4 QUALITATIVE ANALYSIS

4.1 CAFQA benefits on a micro-benchmark

In Fig.5 we use a 2-qubit ‘XX’ Hamiltonian system and a 2-qubit hardware-efficient ansatz with only one tuning parameter to show the benefits of CAFQA. The Y-axis shows the estimated expectation values of the Hamiltonian while the X-axis sweeps the tuning parameter:

① The green line represents tuning the one ansatz parameter on an ideal noise-free quantum device. Sweeping through all rotations produces an expectation value minima = -1.0.

② Next, the same tuning is performed on two noisy quantum devices, IBMQ Casablanca and Manhattan (simulated with noise models). These are shown with the purple and blue lines. Clearly,




Figure 5: Comparison of different methods of ansatz tuning for a 2qubit XX Hamiltonian employing an ansatz with only one tunable rotation angle parameter. Although the search space for CAFQA is limited, it is able to achieve the global minimum (equaling the ideal machine). Furthermore, CAFQA outperforms the noisy machines, which are limited by noise although they can explore the entire tuning space. On the other hand, the HF method of initialization is unable to recover the expectation value at all since the XX Hamiltonian does not have an uncorrelated component.

the noisy devices are able to sweep through the entire parameter space, but the effect of noise limits the minimum obtained, achieving only -0.7 / -0.85. Note that this microbenchmark is too simplistic to suffer from barren plateaus, etc., but it is expected that the deviation

from the ideal / exact minimum will increase with more complex problems and increasing noise.

③ Next, the dashed line shows the expectation value produced by HF initialization. In this example, HF is unable to produce any useful result since the chosen Hamiltonian does not have any diagonal Pauli terms suitable for HF. This can be thought to represent pure electron correlation energy in the context of molecular chemistry (as described in Section 2.2).

④ Finally, CAFQA is shown with orange 'X's. Note that there are only 4 unique discrete points in the Clifford space for the one tunable parameter. Even so, CAFQA is able to produce the expectation value global minimum ($= -1.0$) via one of those 4 Clifford points. Not only does the Clifford minimum match the minimum of the entire tuning space, but also the ability to simulate the Clifford space ideally without noise produces more accurate estimations than the noisy devices.

4.2 High accuracy CAFQA stabilizer states

In Fig.6 we break down the expectation value returned by the Clifford Ansatz in comparison to Hartree-Fock and ideal / exact minimum from noise free simulation. This is shown for the ground-state energy estimation of the LiH molecule, represented by a 4 qubit Hamiltonian system, at a bond length of 4.8Å. The Y-axis shows the expectation value of each Pauli term and the X-axis lists the Pauli terms in the Hamiltonian.

① In the figure, the resulting expectation value of each Pauli term for the Hartree-Fock (HF) initialization is shown in blue. Since HF is a 'classical' computational basis state (i.e., a bitstring), and therefore Clifford, as noted in Section 3, all the HF Pauli term expectation values are $+1 / -1 / 0$. Further, since HF is 'classical', all non-diagonal Pauli terms (i.e., any terms apart from the tensor products of I and Z) have an expectation value of zero. Only calculating the expectation values for the diagonal terms leads to HF ignoring the correlation energy, which is known to cause serious errors for some larger molecules (described in Section 2.2).

② The expectation value of each Pauli term for the Clifford ansatz produced by CAFQA is shown in red. Again, note that all expectation values are $+1 / -1 / 0$. Moreover, note that for the Clifford ansatz, there are multiple non-diagonal Pauli terms which produce an expectation of $+1 / -1$. The non-zero expectation on non-diagonal terms is indicative of CAFQA producing a non-computational basis state, albeit a Clifford one. By doing so, it is able to capture some of the correlation energy that is contributed by the non-diagonal Pauli terms. This is important because it is qualitatively indicative of the potential for high(er) estimation accuracy through CAFQA as the complexity of the problem scales.

③ The expectation value of each Pauli term for the exact minimum from ideal noise-free simulation is shown in green. Ideal noise-free simulation is possible since LiH is a very small molecule. While expectation values range from -1 to 1 , it is evident that the expectation values are close to those produced by the Clifford Ansatz, both for diagonal terms and the non-diagonal terms. This is a clear indicator of the effectiveness of CAFQA for ground-state energy estimation for the LiH molecule (at the chosen bond length). This is confirmed later in Section 7.1.2 which shows that CAFQA is able to achieve high accuracy in the range of 10^{-2} Hartree for

LiH. This trend is also observed across other molecules and bond lengths, leading to high accuracy overall.

5 DISCRETE SEARCH OVER THE CLIFFORD SPACE

To efficiently search through the discrete Clifford parameter space, CAFQA requires a sample-efficient search technique to find the performant ansatz parameters as quickly as possible. Bayesian optimization is one such technique that actively and intelligently queries the most informative samples at each round to reduce the number of samples required [23]. Due to this merit, Bayesian optimization has been successfully applied to different domains in computer systems optimization such as compiler tuning [50], resource allocation [53], and configuration optimization [6, 18, 19, 62]. Bayesian optimization iteratively alternates between intelligent sampling and model updates. As such, it includes two components: a surrogate model and an acquisition function. The surrogate model tries to learn the unknown underlying function that maps the search parameters to the problem objective (e.g., ground-state energy). The acquisition function is the search strategy that selects the next sample to query to update the surrogate model.

CAFQA searches through the Clifford space with Bayesian optimization to identify optimal Clifford gates for the tunable circuit parameters. Each tunable parameter is able to take one of four rotational angles as was shown in Fig.3. This creates a discrete search space complexity of $O(4^{\#params})$ that scales exponentially in the number of parameters, although it is considerably smaller than the entire quantum tuning space. While Bayesian Optimization efficiency can degrade with increased number of search parameters (or dimensionality) [40], it is still observed to be effective in searching through the Clifford space since each parameter chooses only from four different rotational angles.

Since the Clifford parameter space is discrete, CAFQA chooses the random forest as the surrogate model as it is flexible enough to model the discrete space and scales well [50]. CAFQA uses a greedy acquisition function [18, 50] to select samples with the lowest energy estimates predicted from the surrogate model. Empirically, the combination of the random forest surrogate model and the greedy acquisition function gives highly accurate results, as illustrated in Section 7. Details on the implementation of the search algorithm can be found in [50]. Here, we limit ourselves to an illustrative example.

Fig.7 shows the discrete search employed by CAFQA to produce a Clifford Ansatz for H₂O ground state energy estimation at a bond length of 4Å. The first 1,000 iterations are a warm-up period, which involves randomly sampling and mapping the search space, a key component to BO. The search algorithm then uses these random samples to efficiently search the parameter space. In the figure, note that as soon as the random sampling is complete, the search algorithm begins to find better expectation values compared to random. A potentially global minimum is found after an additional 600 search iterations. Notably, in this use case, the identified minimum is well within the chemical accuracy requirements. Although 2000 iterations are shown here, the search can be constrained by a tuning budget or by the saturation of the obtained minimum.




Figure 6: LiH ground state energy at bond length of 4.8 Å (3x equilibrium). Expectation value of each Pauli term, as obtained from different methods, is shown. The Pauli terms along the X-axis are arranged as: i) Computational basis terms, ii) Non computational basis terms selected by CAFQA, and iii) remaining terms which are beyond the Clifford reach (sorted by Exact expectation value). While HF is only able to obtain non zero expectations (of ± 1) for diagonal Pauli (computational basis) terms, the Clifford Ansatz enables expectations of ± 1 for non-diagonal Pauli terms as well. Further, the similarity in expectation value between the Clifford Ansatz and the exact (i.e., ideal) LiH tuning is evident.




Figure 7: H₂O ground state energy estimation via CAFQA’s Clifford ansatz discrete search at bond length of 4 Å (4x equilibrium). The first 1000 iterations are a warm up period. Beyond this, the search can achieve estimations within the chemical accuracy target in an additional 600 iterations. In this instance, post-CAFQA variational tuning on a quantum machine is not required.

App	# Qu.	Bond Len. (Eqbm.)	Bond Len. (Range)	Mol Orbitals Total / Used
H ₂	2	0.74 Å	0.37 - 2.96 Å	2 / 2
LiH	4	1.6 Å	0.8 - 4.8 Å	4 / 3
H ₂ O	12	1 Å	0.5 - 4.0 Å	7 / 7
H ₆	10	0.9 Å	0.45 - 3.6 Å	6 / 6
N ₂	12	1.09 Å	0.55 - 4.36 Å	10 / 7
Cr ₂	34	1.68 Å	1.25 - 3.5 Å	36 / 18
NaH	12	1.9 Å	0.95 - 7.6 Å	10 / 7
H ₂ -S1	18	-	-	-
BeH ₂	12	1.32 Å	0.66 - 5.28 Å	7 / 7

Table 1: VQA applications and their characteristics.

6 METHODOLOGY

Ground state energy estimation of molecules: We use VQE to estimate the ground state energy of the following molecules: H₂, LiH, H₂O, H₆, Cr₂, N₂, NaH, H₂-S1 and BeH₂. Hamiltonians are constructed in the STO-3G basis with parity mapping and Z2 symmetry / two qubit reduction. Hamiltonians are constructed for spin corresponding to the singlet (0 unpaired electrons in the orbitals) electronic state, which usually has the lowest energy near equilibrium geometries (more on this in Section 7.1.4). We provide detailed evaluations for the first five and only mean accuracy results for the other three. H₂-S1_STO-3G_singlet (H₂-S1) is obtained from Contextual Subspace VQE [37]. Details about these molecules and their representative Hamiltonians are provided in Table 1.

H₆ is known to be a prototypical strongly correlated molecule, thus widening the gap between ideal results and classical methods. Also notable is the Chromium Dimer (Cr₂), which has long been a benchmark molecule for evaluating the performance of different computational methods due to its unusual bonding properties in its ground and excited states [21, 76]. Cr₂ is especially challenging to simulate, requiring a system of as many as 72 qubits; therefore, we are unable to compare against its exact estimates. Furthermore, we freeze the lower 18 (out of Cr₂’s 36 orbitals) to reduce the system to 34 qubits, to ease the burden of iterative tuning given our reasonable yet limited computational resources - but this is not a strict limitation. Freezing lower orbitals is least detrimental to bond dissociation energy estimations - electrons closer to the nucleus are tightly attached and have high ionization energies [4].

For all Hamiltonians above, we use a hardware-efficient SU2 parameterized circuit [1] with one layer of linear entanglement as ansatz. An example of this for 10 qubits is shown in Fig. 3. Different initialization comparisons are performed on this circuit.

Evaluation Comparisons: We compare the following -

① CAFQA: Our proposed approach, which uses a Clifford-only ansatz, and potentially produces the best possible stabilizer initial state for the target Hamiltonian.

② *Exact*: The exact energy estimations computed classically (but possible only for small problem sizes).

③ *Hartree-Fock (HF)*: HF is the best computational basis state for the target Hamiltonian under specified electron and spin preservation constraints.

Evaluation Metrics: We evaluate CAFQA across four metrics detailed below:

① *Ground State Energy*: Potential energy as a function of nuclear coordinates, as estimated by different techniques, expressed in Hartree units.

② *Energy estimation accuracy*: Absolute energy difference between energy estimates from different techniques and exact estimates, expressed in Hartree units. Chemical accuracy region is shown in orange.

③ *Recovered correlation energy*: Percentage of the difference between the Exact energy and the Hartree-Fock limit that is recovered by CAFQA.

④ *Relative accuracy*: Relative energy estimation accuracy between CAFQA and state-of-the-art HF (only Fig.13).

Infrastructure: The CAFQA framework is implemented in Python. Hartree-Fock estimations are performed via Psi4 [74] while CAFQA evaluations are performed using Qiskit [5]. Qiskit interfaces with the PySCF library [68] in the process of constructing Hamiltonians from molecular specifications. The discrete search to find the optimal Clifford gates is performed through Bayesian Optimization via the HyperMapper [50] framework. Classical computations are predominantly carried out on the Google Compute Cloud.

7 EVALUATION

7.1 Detailed molecular analysis

In Figures 8-11 we show ground state energy estimation through VQE for 4 molecules, over different bond lengths. Each figure shows different evaluation metrics for the target molecule: the top subfigure shows the absolute ground state energy (in Hartree), the middle subfigure shows the error in energy estimation, and the bottom subfigure shows the correlation energy recovered by CAFQA over HF. We compare CAFQA in green against exact evaluations in orange; and Hartree-Fock initialization (HF) in blue. More details on these molecules, metrics and comparisons are discussed in Section 6.

7.1.1 H_2 . First, we look at H_2 shown in Fig.8 (a)-(c). In Fig.8 (a) and (b), we see that the HF steadily deviates away from exact energies as bond lengths increase. This is not surprising, as HF is known to work best at / near equilibrium geometry, as discussed in Section 2.2. CAFQA matches HF at low bond lengths but achieves lower energy estimates at higher bond lengths, thus being closer to exact estimates.

Fig.8 (a) also shows CAFQA energy estimates for the H_2^+ cation. The cation is in a higher energy state than its neutral counterpart; this is intuitive, as H_2 does not naturally ionize. For a given molecular system, the Fock (i.e., energy) space represented by the problem Hamiltonian combines the energy spaces of the molecular forms with all possible numbers of electrons and all electron-spin combinations [63]. Thus, when solving for the ground-state energy of




Figure 8: Dissociation curves for H_2 . Evaluation of CAFQA in terms of ground state energy, energy estimation error and correlation energy recovered. Comparisons to Exact / Chemical Accuracy and Hartree-Fock are shown. Energy for H_2^+ cation is also shown.

higher-energy cations and anions, some explicit enforcement of electron preservation constraints on VQE is often required (whereas this is usually not required for the lowest-energy neutral molecule). For H_2^+ , this means that only electronic energies corresponding to a one electron system should be considered for VQE. With CAFQA, we impose electron count constraints through the search objective function. Thus, CAFQA ensures that any required constraints are maintained, along the lines of prior work [63].

Fig.8 (b) shows that CAFQA's error is always less than 10^{-1} Hartree, and is able to achieve estimates near chemical accuracy. On the other hand, as the bond lengths increase, the HF error is 10^{-1} Hartree and greater. Thus, the benefits of CAFQA are clearly evident. Finally, Fig.8 (c) shows that CAFQA is able to recover up to 99.7% of the correlation energy as bond lengths increase.

Overall, CAFQA achieves more accurate energy estimates than HF's best computational basis state and is able to incorporate the expectation of non-diagonal Pauli terms (as discussed in Section 4.2). It produces a non-computational basis state which is uncommon in other classical approaches - an example of this was illustrated earlier for LiH in Fig.6. Furthermore, the stabilizer state produced appears intuitively suited to a molecule such as H_2 . The opposing attractive and repulsive forces of similar strengths acting on the electrons (especially at higher bond lengths) can result in optimal configurations bearing resemblance to stabilizers. Further examination at a molecular level is beyond our current scope.

7.1.2 LiH. Next, we examine LiH shown in Fig.9 (a)-(c). In Fig.9 (a) and (b), we see that HF deviates considerably from the exact value at medium-high bond lengths, but is closer to the exact value at low bond lengths. As before, CAFQA is closest to exact, especially accurate at low and high bond lengths, but always achieves equal or more accurate energy estimates compared to HF. Fig.9 (b) shows that CAFQA's estimation error is usually in the 10^{-1} to 10^{-2} Hartree range, with higher accuracy at high bond lengths. Finally, Fig.9 (c) shows that CAFQA is able to recover up to 93% of the correlation energy at medium-high bond lengths. As before, improving beyond HF, CAFQA is able to produce a non-computational basis state as its ansatz initialization state.

7.1.3 H_2O . Next, we look at H_2O shown in Fig.10 (a)-(c). The first aspect to be noted is that the HF Psi4 estimations do not converge at high bond lengths, so we extrapolate the expected trend as shown in Fig.10 (a). In Fig.10 (a) and (b), we see that HF steadily deviates away from the exact at higher bond lengths. CAFQA matches HF at lower bond lengths but achieves considerably better energy estimates compared to HF at medium / higher bond lengths.

It is interesting to observe the kink in the energy estimation near a bond length of 1.5 Å. This appears to match the prior observation that this is caused by the energy crossing of the lowest singlet (0 unpaired electrons) and triplet (2 unpaired electrons) electronic states for the H_2O molecule [63]. Accordingly, the singlet and triplet states from CAFQA are plotted in pink and yellow, respectively.

Next, Fig.10 (b) shows that CAFQA is able to impressively achieve chemical accuracy at higher bond lengths while HF has high error in the range of 10^{-1} Hartree. At lower bond lengths, CAFQA achieves error rates of around 10^{-2} Hartree. Finally, Fig.10 (c) shows that CAFQA is able to recover up to 99.998% of the correlation energy over HF as the bond lengths increase.




Figure 9: Dissociation curves for LiH. Evaluation of CAFQA in terms of ground state energy, energy estimation error and correlation energy recovered. Comparisons to Exact / Chemical Accuracy and Hartree-Fock are shown.




(a) H_2O Energy(b) H_2O Accuracy(c) H_2O Correlation Energy

Figure 10: Dissociation curves for H_2O . Evaluation of CAFQA in terms of ground state energy, energy estimation error and correlation energy recovered. Comparisons to Exact / Chemical Accuracy and Hartree-Fock are shown. CAFQA (s) and (t) refer to energies corresponding to singlet and triplet states.




(a) H_6 Energy(b) H_6 Accuracy(c) H_6 Correlation Energy

Figure 11: Dissociation curves for H_6 . Evaluation of CAFQA in terms of ground state energy, energy estimation error and correlation energy recovered. Comparisons to Exact / Chemical Accuracy and Hartree-Fock are shown. CAFQA opt. refers to best estimates over multiple spin-optimized Hamiltonians.




Figure 12: Ground state energy for Cr_2 with lower 18 out of 36 orbitals frozen. CAFQA’s estimates are limited by computational time / resources, thus Clifford estimates can be suboptimal. Comparison to HF is shown.

7.1.4 H_6 . Next, we look at H_6 shown in Fig. 11 (a)-(c). In Fig. 11 (a) and (b) we show two versions of HF and CAFQA. The traditional HF and CAFQA results correspond to Hamiltonians generated for orbitals optimized for the singlet electronic state as mentioned in Section 6. On the other hand, the optimized results (labeled ‘opt.’) are produced by generating unique Hamiltonians for different spins and with orbitals optimized accordingly. Then, the HF and CAFQA results corresponding to the lowest estimates across all Hamiltonians are selected for every bond length. It is evident that ‘opt.’ results produce better energy estimates compared to those obtained from the singlet-optimized Hamiltonian at higher bond lengths. This shows that optimizing the Hamiltonian to the best extent / as widely as possible can considerably improve VQE estimation, at the cost of increased compute. The HF techniques and CAFQA are far from exact while CAFQA opt. is near exact at high bond lengths. Deviations from the exact value are not surprising, as H_6 has high correlation energy.

Fig. 11 (b) shows that CAFQA errors are in the 10^{-1} Hartree range (except at high bond lengths when CAFQA opt. achieves the chemical accuracy range), thus clearly requiring quantum exploration on a quantum device post-ansatz selection. Finally, Fig. 11 (c) shows that CAFQA is able to recover up to 50% of the correlation energy over HF as the bond lengths increase, while CAFQA opt. can achieve near 100% at high bond lengths. CAFQA is again able to produce a non-computational basis state as its ansatz initialization state, although it is evident that exploration of only the Clifford space limits accuracy.

7.1.5 Cr_2 . In Fig. 12 we evaluate CAFQA on Cr_2 for ground state energy estimation. For Cr_2 , we are unable to generate exact evaluations, since the size of the system is too large (discussed in Section 6). We compare CAFQA to HF. It is evident from the figure that CAFQA consistently achieves better initial energy estimates compared to HF across all bond lengths. In addition, CAFQA has resemblance to experimental estimates [76], although there are some limitations to the comparison due to the orbital freezing that we utilize (discussed




Figure 13: CAFQA accuracy compared to state-of-the-art Hartree-Fock. ‘Average’: Relative error reduction averaged over all bond lengths for each molecule. ‘Maximum’: Maximum error reduction for each molecule, usually at the greatest bond length.

in Section 6), as well as a lack of sufficient molecular specifications. We note that very recent work [38] (subsequent to CAFQA) showed high accuracy computational prediction of the Cr_2 potential energy curve, consistent with experimental data. Comparisons to this work are worth pursuing.

Due to the size of the problem, we are limited by resources in running the Bayesian Optimization search extensively at each bond length. Limited search means that CAFQA at some bond lengths can produce sub-optimal estimates. Our current estimates are obtained over a 1-week period but these estimations can improve with more memory / compute, more execution time, better search strategies, efficient parallelization, limited exploration of non-Cliffords — more in Section 8. We reemphasize that CAFQA is only a first step in VQA tasks, with the primary goal of producing an ansatz initial state well suited to further quantum exploration on a quantum device.

7.2 Relative accuracy compared to SOTA

Fig. 13 shows the accuracy achieved by CAFQA in all applications (except Cr_2) - 8 VQE molecular chemistry ground state estimation tasks, relative to the state-of-the-art Hartree-Fock approach. Two sets of results are shown: ‘Average’ and ‘Maximum’. For ‘Average’, the relative error reduction of CAFQA compared to HF is averaged across all the evaluated bond lengths (for each molecule). For ‘Maximum’, the highest error reduction of CAFQA compared to HF is presented, which is usually at the greatest bond length, since Hartree-Fock steadily deteriorates away from equilibrium.

It is evident that the CAFQA is able to achieve significant average relative accuracy improvements over all applications, with a mean of 6.4x (highest of 25x). Furthermore, the maximum improvements are very substantial, with a mean of 56.8x (highest of 3.4×10^5 x). The lowest benefits are obtained for H_6 , which, as explained in Section 6, has a significant correlation energy component which cannot be entirely recovered by only exploring the Clifford space. It is clearly evident that high VQA initialization accuracy can be achieved by CAFQA compared to state-of-the-art.

7.3 Post-CAFQA VQA exploration

Although today’s NISQ machines are often too noisy to improve on CAFQA’s estimates, it is expected that NISQ machines in the near future will be able to do so. In this case, the CAFQA initialization will allow for more focused tuning on the machine, resulting in lower




Figure 14: Post-CAFQA VQA tuning for LiH. CAFQA initialization leads to 2.5x faster VQA convergence compared to HF, on both ideal and noisy machines.




Figure 15: BO search iterations for CAFQA to converge to the lowest expectation estimates for each VQA problem. Search overheads are reasonable, especially considering the benefits of reduced variational tuning on NISQ devices.

potential for detrimental impact from noise and barren plateaus, thereby leading to faster and more accurate convergence.

This is illustrated in Fig. 14 which shows post-CAFQA VQE tuning for LiH ground state energy estimation. Evaluation is shown for tuning beginning from HF initialization and CAFQA initialization, respectively. Furthermore, two sets of results are shown, one on ideal noise-free simulation and the other on noisy simulation modeled on real machine characteristics. In both sets of results it is evident that CAFQA-initialized exploration converges roughly 2.5x faster than HF-initialization, clearly indicative of the benefits from better initialization.

It can be observed that the ideal simulation produces near exact results, improving over the initialization. Furthermore, the energy estimate produced by noisy simulation (error roughly = 10^{-2} Hartree) is on par with the estimate obtained directly from CAFQA initialization itself. While we do not expect the latter trend to hold for more complex Hamiltonians and as machine noise reduces, CAFQA initialization will continue to be useful for fast and accurate convergence. Greater benefits can be expected for larger problem sizes, which can be realistically evaluated as NISQ machines improve. Reduced execution on the actual quantum device is also beneficial from the monetary standpoint. Prior work discusses high execution costs of variational algorithms on the quantum cloud, consuming thousands of dollars to execute problems of reasonably small sizes [31].

7.4 Discrete Search

Fig. 15 shows the number of iterations consumed by CAFQA’s discrete search to converge to a minimum energy estimate. It is evident that the number of iterations increases with the size of the problem because the number of tuning parameters increases. The number of iterations across all applications is very reasonable considering the benefits of reduced variational tuning on noisy quantum devices. The current run time of CAFQA varies roughly from a few minutes (H_2) to a week (Cr_2). The execution time can be reduced via increased compute / memory resources, improved search algorithm, parallel search, etc.

8 DISCUSSION

Simulation beyond Cliffords: Prior work has shown that efficient classical simulation can be extended beyond Clifford-only circuits to constrained Clifford+T circuits wherein T refers to the single-qubit 45-degree phase shift [11, 12]. Optimally designing a CAFQA ansatz with a mix of Clifford gates and minimal T gates is worth exploring. We perform preliminary exploration of allowing a few T gates within the CAFQA framework. Note that the simulation complexity grows exponentially with the number of T gates, so the number and location of the T gates require careful analysis. Our current exploration only studies the insertion of T gates at prior Clifford gate positions in the ansatz and only attempts the addition of under 10 T gates. But it is already evident that this direction is promising — Fig. 16 shows that the addition of just up to 1 T gate for H_2 and up to 4 T gates for LiH significantly improves initialization accuracy, while remaining classically simulable. $CAFQA+kT$ ($k \leq 1 / k \leq 4$) is able to recover as much as 99.9% of the correlation energy at bond lengths for which Clifford-only CAFQA accuracy is relatively limited.

Beyond a hardware-efficient ansatz: A hardware-efficient ansatz can be limited in its capabilities because it is application-agnostic. Thus, expanding beyond this ansatz can be beneficial if efficiently suited to the CAFQA approach. Expanding the ansatz search to Clifford plus limited non-Clifford gates (discussed above) is a first step in this direction - potentially allowing for a dynamically evolving ansatz structure similar to ADAPT-VQE [29].

Optimization: The discrete search via Bayesian Optimization employed by CAFQA is able to produce high accuracy results and in a reasonable number of iterations and runtime. Although this is clearly efficient for the target problem space (our largest is a 34-qubit system), the search could face scalability challenges on significantly larger problems. This is especially important since the Clifford search space (stabilizer state space) scales exponentially in the number of tunable parameters (qubits) [33]. Thus, optimizing the search strategy at the algorithmic as well as implementation levels can reap benefits.

Related Work: [56] proposes initializing ansatz parameters such that subsections of the ansatz do not form a ‘2-design’. [45] proposes a perturbative expansion of the cost function from HF initialization. This can produce an initialization state very close to the HF, but potentially performs better. FLIP [64] proposes initialization with the help of machine learning. MetaVQE [14] encodes the Hamiltonian parameters in the first layers of the quantum circuit.




Figure 16: CAFQA + kT dissociation curves, for H_2 and LiH . The addition of just up to 1 T gate for H_2 and up to 4 T gates for LiH is seen to significantly improve initialization, while remaining classically simulable.

9 CONCLUSION

Advancing NISQ frontiers to real world applicability requires concerted effort on multiple fronts, with support from sophisticated error mitigation, classical computing, and more. In this spirit, CAFQA proposes application-specific classical simulation bootstrapping for VQAs. CAFQA tackles the problem of finding initial VQA parameters by proposing a “Clifford Ansatz” — an ansatz that is a hardware efficient circuit built with only Clifford gates. In this ansatz, the initial parameters for the tunable gates are chosen by searching efficiently through the Clifford parameter space via classical simulation, and thereby producing a suitable high accuracy initial state that outperforms state-of-the-art approaches. Furthermore, there is considerable potential to extend these findings beyond Cliffords and to other circuit structures.

CAFQA is a promising example of quantum-inspired classical techniques as a supporting methodology for VQAs in the NISQ era and beyond. It also highlights the potential for a synergistic quantum-classical paradigm to boost NISQ-era quantum computing towards real world applicability.

ACKNOWLEDGMENTS

This work is funded in part by EPIQC, an NSF Expedition in Computing, under award CCF-1730449; in part by STAQ under award NSF Phy-1818914; in part by NSF award 2110860; in part by the US Department of Energy Office of Advanced Scientific Computing Research, Accelerated Research for Quantum Computing Program; and in part by the NSF Quantum Leap Challenge Institute for Hybrid Quantum Architectures and Networks (NSF Award 2016136) and in part based upon work supported by the U.S. Department of Energy, Office of Science, National Quantum Information Science Research Centers. GSR is supported as a Computing Innovation Fellow at the University of Chicago. This material is based upon work supported by the National Science Foundation under Grant # 2030859 to the Computing Research Association for the CIFellows Project. KNS is supported by IBM as a Postdoctoral Scholar at the University of Chicago and the Chicago Quantum Exchange. YD’s work is supported by the National Science Foundation under Grant # 2030859 to the Computing Research Association for the CIFellows Project. HH is supported by NSF (grants CCF-2119184, CNS-1956180, CNS-1952050, CCF-1823032, CNS-1764039), ARO (grant W911NF1920321), and a DOE Early Career Award (grant DESC0014195 0003). WMK acknowledges funding from NSF Grant DGE-1842474. FTC is Chief Scientist for Quantum Software at ColdQuanta and an advisor to Quantum Circuits, Inc.

REFERENCES

- [1] 2021. IBM Quantum SU2 ansatz. <https://qiskit.org/documentation/stubs/qiskit.circuit.library.EfficientSU2.html>.
- [2] 2021. Microsoft Docs Hartree–Fock Theory. <https://docs.microsoft.com/en-us/azure/quantum/user-guide/libraries/chemistry/concepts/hartree-fock>.
- [3] 2021. Techniques and Applications of Quantum Monte Carlo. <https://web.ornl.gov/~kentpr/thesis/pkthnode13.html>.
- [4] 2022. NIST Atomic Spectra Database Ionization Energies Form. <https://physics.nist.gov/PhysRefData/ASD/ionEnergy.html>. Accessed: 2022-01-01.
- [5] Héctor Abraham, AduOfiei, Rochisha Agarwal, Ismail Yunus Akhalwaya, Gadi Aleksandrowicz, Thomas Alexander, Matthew Amy, Eli Arbel, Arijit02, Abraham Asfaw, Artur Avkhadiev, Carlos Azaustre, AzizNgoueya, Abhik Banerjee, Aman Bansal, Panagiotis Barkoutsos, George Barron, George S. Barron, Luciano Bello, Yael Ben-Haim, Daniel Bevenius, Arjun Bhohe, Lev S. Bishop, Carsten Blank, Sorin Bolos, Samuel Bosch, Brandon, Sergey Bravyi, Bryce-Fuller, David Bucher, Artemiy Burov, Fran Cabrera, Padraic Calpin, Lauren Capelluto, Jorge Carballo, Ginés Carrascal, Adrian Chen, Chun-Fu Chen, Edward Chen, Jiellun (Chris) Chen, Richard Chen, Jerry M. Chow, Spencer Churchill, Christian Claus, Christian Clauss, Romilly Cocking, Filipe Correa, Abigail J. Cross, Andrew W. Cross, Simon Cross, Juan Cruz-Benito, Chris Culver, Antonio D. Córcoles-Gonzales, Sean Dague, Tareq El Dandachi, Marcus Daniels, Matthieu Dartiaillh, DavideFrr, Abdón Rodríguez Davila, Anton Dekusar, Delton Ding, Jun Doi, Eric Drechsler, Drew, Eugene Dumitrescu, Karel Dumon, Ivan Duran, Kareem EL-Safty, Eric Eastman, Grant Eberle, Pieter Eendebak, Daniel Egger, Mark Everitt, Paco Martín Fernández, Axel Hernández Ferrera, Romain Fouilland, FranckChevallier, Albert Frisch, Andreas Fuhrer, Bryce Fuller, MELVIN GEORGE, Julien Gacon, Borja Godoy Gago, Claudio Gambella, Jay M. Gambetta, Adhisha Gammanpila, Luis Garcia, Tanya Garg, Shelly Garion, Austin Gilliam, Aditya Giridharan, Juan Gomez-Mosquera, Salvador de la Puente González, Jesse Gorzinski, Ian Gould, Donny Greenberg, Dmitry Grinko, Wen Guan, John A. Gunnels, Mikael Haglund, Isabel Haide, Ikko Hamamura, Omar Costa Hamido, Frank Harkins, Vojtech Havlicek, Joe Hellmers, Łukasz Herok, Stefan Hillmich, Hiroshi Horii, Connor Howington, Shaohan Hu, Wei Hu, Junye Huang, Rolf Huisman, Haruki Imai, Takashi Imamichi, Kazuaki Ishizaki, Raban Iten, Toshinari Itoko, JamesSeaward, Ali Javadi, Ali Javadi-Abhari, Jessica, Madhav Jivrajani, Kiran Johns, Scott Johnston, Jonathan-Shoemaker, Vismay K, Tal Kachmann, Naoki Kanazawa, Kang-Bae, Anton Karazeev, Paul Kassebaum, Josh Kelso, Spencer King, Knabberjoe, Yuri Kobayashi, Arseny Kovyrshin, Rajiv Krishnakumar, Vivek Krishnan, Kevin Krsulich, Prasad Kumkar, Gaweil Kus, Ryan LaRose, Enrique Lacal, Raphaël Lambert, John Lapeyre, Joe Latone, Scott Lawrence, Christina Lee, Gushu Li, Dennis Liu, Peng Liu, Yunho Maeng, Kahan Majmudar, Aleksei Malyshev, Joshua Manela, Jakub Marecek, Manoel Marques, Dmitri Maslov, David Mathews, Atsushi Matsuo, Douglas T. McClure, Cameron

- McGarry, David McKay, Dan McPherson, Srujan Meesala, Thomas Metcalfe, Martin Mevissen, Andrew Meyer, Antonio Mezzacapo, Rohit Midha, Zlatko Minev, Abby Mitchell, Nikolaj Moll, Jhon Montanez, Michael Duane Mooring, Renier Morales, Niall Moran, Mario Motta, MrF, Prakash Murali, Jan Müggenburg, David Nadlinger, Ken Nakanishi, Giacomo Nannicini, Paul Nation, Edwin Navarro, Yehuda Naveh, Scott Wyman Neagle, Patrick Neuweiler, Johan Nicander, Pradeep Niroula, Hassi Norlen, NuoWenLei, Lee James O'Riordan, Oluwatobi Ogunbayo, Pauline Ollitrault, Raul Otaolea, Steven Oud, Dan Padilha, Hanhee Paik, Soham Pal, Yuchen Pang, Simone Perriello, Anna Phan, Francesco Piro, Marco Pistoia, Christophe Piveteau, Pierre Pocreau, Alejandro Pozas-Kerstjens, Viktor Prutyay, Daniel Puzzuoli, Jesús Pérez, Quintiii, Rafey Iqbal Rahman, Arun Raja, Nipun Ramagiri, Anirudh Rao, Rudy Raymond, Rafael Martin-Cuevas Redondo, Max Reuter, Julia Rice, Marcello La Rocca, Diego M. Rodriguez, RohithKarur, Max Rossmannek, Mingi Ryu, Tharrmashastha SAPV, SamFerracin, Martin Sandberg, Hirnay Sandesara, Ritvik Sapra, Hayk Sargsyan, Aniruddha Sarkar, Ninad Sathaye, Bruno Schmitt, Chris Schnabel, Zachary Schoenfeld, Travis L. Scholten, Eddie Schoute, Joachim Schwarm, Ismael Faro Sertage, Kanav Setia, Nathan Shammah, Yuhong Shi, Adenilton Silva, Andrea Simonetto, Nick Singstock, Yukio Siraichi, Iskandar Sitdikov, Seyon Sivarajah, Magnus Berg Sletfjerd, John A. Smolin, Mathias Soeken, Igor Olegovich Sokolov, Igor Sokolov, SooluThomas, Starfish, Dominik Steenken, Matt Stypulkoski, Shaojun Sun, Kevin J. Sung, Hitomi Takahashi, Tanvish Takawale, Ivano Tavernelli, Charles Taylor, Pete T aylour, Soolu Thomas, Mathieu Tillet, Maddy Tod, Miroslav Tomasik, Enrique de la Torre, Kenso Traming, Matthew Treinish, TrishaPe, Davindra Tulsi, Wes Turner, Yotam Vaknin, Carmen Recio Valcarce, Francois Varchon, Almudena Carrera Vazquez, Victor Villar, Desiree Vogt-Lee, Christophe Vuillot, James Weaver, Johannes Weidenfeller, Rafal Wieczorek, Jonathan A. Wildstrom, Erick Winston, Jack J. Woehr, Stefan Woerner, Ryan Woo, Christopher J. Wood, Ryan Wood, Stephen Wood, Steve Wood, James Wootton, Daniyar Yeralin, David Yonge-Mallo, Richard Young, Jessie Yu, Christopher Zachow, Laura Zdanski, Helena Zhang, Christa Zoufal, Zoufale, a kapila, a matsuo, becamorrison, brandhsn, nick bronn, chlorophyll zz, dekelmeirom, dekelmeirom, dekol, dime10, drholmie, dtrenev, echen, elfrocampador, faisaldebouni, fanizzamarco, gabrieleagl, gadial, galeinston, georgios ts, gruu, hhorii, hykavitha, jagunther, jliu45, jscott2, kanejess, klinvill, krutic2966, kurarr, lerongil, ma5x, merav aharoni, michelle4654, ordmoj, sagar pahwa, rmo-yard, saswati qiskit, scottkelso, strickroman, sumitpuri, tigerjack, toural, tsura crisaldo, vvlpas, welien, willhbang, yang.luh, yotamvakninibm, and Mantas Cepulkovskis. 2019. Qiskit: An Open-source Framework for Quantum Computing. <https://doi.org/10.5281/zenodo.2562110>
- [6] Omid Alipourfard, Hongqiang Harry Liu, Jianshu Chen, Shivaram Venkataraman, Minlan Yu, and Ming Zhang. 2017. Cherrypick: Adaptively Unearthing the Best Cloud Configurations for Big Data Analytics. In *Proceedings of the 14th USENIX Conference on Networked Systems Design and Implementation* (Boston, MA, USA) (NSDI'17). USENIX Association, USA, 469–482.
- [7] George S. Barron and Christopher J. Wood. 2020. Measurement Error Mitigation for Variational Quantum Algorithms. arXiv:2010.08520 [quant-ph]
- [8] Jacob Biamonte, Peter Wittek, Nicola Pancotti, Patrick Rebentrost, Nathan Wiebe, and Seth Lloyd. 2017. Quantum machine learning. *Nature* 549, 7671 (2017), 195–202.
- [9] Sergio Boixo, Sergei V. Isakov, Vadim N. Smelyanskiy, Ryan Babbush, Nan Ding, Zhang Jiang, Michael J. Bremner, John M. Martinis, and Hartmut Neven. 2018. Characterizing quantum supremacy in near-term devices. *Nature Physics* 14, 6 (Apr 2018), 595–600. <https://doi.org/10.1038/s41567-018-0124-x>
- [10] Ludmila Botelho, Adam Glos, Akash Kundu, Jaroslaw Adam Mszczak, Özlem Salehi, and Zoltán Zimborás. 2021. Error mitigation for variational quantum algorithms through mid-circuit measurements. arXiv:2108.10927 [quant-ph]
- [11] Sergey Bravyi, Dan Browne, Padraic Calpin, Earl Campbell, David Gosset, and Mark Howard. 2019. Simulation of quantum circuits by low-rank stabilizer decompositions. *Quantum* 3 (Sep 2019), 181. <https://doi.org/10.22331/q-2019-09-02-181>
- [12] Sergey Bravyi and David Gosset. 2016. Improved Classical Simulation of Quantum Circuits Dominated by Clifford Gates. *Physical Review Letters* 116, 25 (Jun 2016). <https://doi.org/10.1103/physrevlett.116.250501>
- [13] M. Cerezo, Akira Sone, Tyler Volkoff, Lukasz Cincio, and Patrick J. Coles. 2021. Cost function dependent barren plateaus in shallow parametrized quantum circuits. *Nature Communications* 12, 1 (Mar 2021). <https://doi.org/10.1038/s41467-021-21728-w>
- [14] Alba Cervera-Lierta, Jakob S. Kottmann, and Alán Aspuru-Guzik. 2021. Meta-Variational Quantum Eigensolver: Learning Energy Profiles of Parametrized Hamiltonians for Quantum Simulation. *PRX Quantum* 2, 2 (May 2021). <https://doi.org/10.1103/prxquantum.2.020329>
- [15] Piotr Czarnik, Andrew Arrasmith, Patrick J. Coles, and Lukasz Cincio. 2020. Error mitigation with Clifford quantum-circuit data. arXiv:2005.10189 [quant-ph]
- [16] Hans De Raedt, Fengping Jin, Dennis Willsch, Madita Willsch, Naoki Yoshioka, Nobuyasu Ito, Shengjun Yuan, and Kristel Michiels. 2019. Massively parallel quantum computer simulator, eleven years later. *Computer Physics Communications* 237 (Apr 2019), 47–61. <https://doi.org/10.1016/j.cpc.2018.11.005>
- [17] Yongshan Ding, Pranav Gokhale, Sophia Fuhui Lin, Richard Rines, Thomas Propson, and Frederic T Chong. 2020. Systematic Crosstalk Mitigation for Superconducting Qubits via Frequency-Aware Compilation. *arXiv preprint arXiv:2008.09503* (2020).
- [18] Yi Ding, Ahsan Pervaiz, Michael Carbin, and Henry Hoffmann. 2021. Generalizable and Interpretable Learning for Configuration Extrapolation. In *Proceedings of the 29th ACM Joint Meeting on European Software Engineering Conference and Symposium on the Foundations of Software Engineering* (Athens, Greece) (ESEC/FSE 2021). Association for Computing Machinery, New York, NY, USA, 728–740. <https://doi.org/10.1145/3468264.3468603>
- [19] Yi Ding, Ahsan Pervaiz, Sanjay Krishnan, and Henry Hoffmann. 2020. Bayesian Learning for Hardware and Software Configuration Co-Optimization. (2020).
- [20] Yuxuan Du, Yang Qian, and Dacheng Tao. 2021. Accelerating variational quantum algorithms with multiple quantum processors. *arXiv preprint arXiv:2106.12819* (2021).
- [21] V. E. Elfving, B. W. Broer, M. Webber, J. Gavartin, M. D. Halls, K. P. Lorton, and A. Bochevarov. 2020. How will quantum computers provide an industrially relevant computational advantage in quantum chemistry? arXiv:2009.12472 [quant-ph]
- [22] Edward Farhi, Jeffrey Goldstone, and Sam Gutmann. 2014. A Quantum Approximate Optimization Algorithm. arXiv:1411.4028 [quant-ph]
- [23] Peter I Frazier. 2018. A tutorial on Bayesian optimization. *arXiv preprint arXiv:1807.02811* (2018).
- [24] Tudor Giurgica-Tiron, Yousef Hindy, Ryan LaRose, Andrea Mari, and William J Zeng. 2020. Digital zero noise extrapolation for quantum error mitigation. In *2020 IEEE International Conference on Quantum Computing and Engineering (QCE)*. IEEE, 306–316.
- [25] Pranav Gokhale, Olivia Angiuli, Yongshan Ding, Kaiwen Gui, Teague Tomesh, Martin Suchara, Margaret Martonosi, and Frederic T Chong. 2019. Minimizing state preparations in variational quantum eigensolver by partitioning into commuting families. *arXiv preprint arXiv:1907.13623* (2019).
- [26] Pranav Gokhale, Yongshan Ding, Thomas Propson, Christopher Winkler, Nelson Leung, Yunong Shi, David I. Schuster, Henry Hoffmann, and Frederic T. Chong. 2019. Partial Compilation of Variational Algorithms for Noisy Intermediate-Scale Quantum Machines. *Proceedings of the 52nd Annual IEEE/ACM International Symposium on Microarchitecture* (Oct 2019). <https://doi.org/10.1145/3352460.3358313>
- [27] Daniel Gottesman. 1998. The Heisenberg representation of quantum computers. *arXiv preprint quant-ph/9807006* (1998).
- [28] Daniel Gottesman and Isaac L Chuang. 1999. Demonstrating the viability of universal quantum computation using teleportation and single-qubit operations. *Nature* 402, 6760 (1999), 390–393.
- [29] Harper R. Grimsley, Sophia E. Economou, Edwin Barnes, and Nicholas J. Mayhall. 2019. An adaptive variational algorithm for exact molecular simulations on a quantum computer. *Nature Communications* 10, 1 (Jul 2019). <https://doi.org/10.1038/s41467-019-10988-2>
- [30] Lov K. Grover. 1996. A Fast Quantum Mechanical Algorithm for Database Search. In *ANNUAL ACM SYMPOSIUM ON THEORY OF COMPUTING*. ACM, 212–219.
- [31] Andi Gu, Angus Lowe, Pavel A Dub, Patrick J Coles, and Andrew Arrasmith. 2021. Adaptive shot allocation for fast convergence in variational quantum algorithms. *arXiv preprint arXiv:2108.10434* (2021).
- [32] Douglas Rayner Hartree and William Hartree. 1935. Self-consistent field, with exchange, for beryllium. *Proceedings of the Royal Society of London. Series A-Mathematical and Physical Sciences* 150, 869 (1935), 9–33.
- [33] Markus Heinrich and David Gross. 2019. Robustness of Magic and Symmetries of the Stabiliser Polytope. *Quantum* 3 (Apr 2019), 132. <https://doi.org/10.22331/q-2019-04-08-132>
- [34] Zoë Holmes, Kunal Sharma, M. Cerezo, and Patrick J. Coles. 2021. Connecting ansatz expressibility to gradient magnitudes and barren plateaus. arXiv:2101.02138 [quant-ph]
- [35] Thomas Häner and Damian S. Steiger. 2017. 0.5 petabyte simulation of a 45-qubit quantum circuit. *Proceedings of the International Conference for High Performance Computing, Networking, Storage and Analysis* (Nov 2017). <https://doi.org/10.1145/3126908.3126947>
- [36] Abhinav Kandala, Antonio Mezzacapo, Kristan Temme, Maika Takita, Markus Brink, Jerry M Chow, and Jay M Gambetta. 2017. Hardware-efficient variational quantum eigensolver for small molecules and quantum magnets. *Nature* 549, 7671 (2017), 242–246.
- [37] William M. Kirby, Andrew Tranter, and Peter J. Love. 2021. Contextual Subspace Variational Quantum Eigensolver. *Quantum* 5 (May 2021), 456. <https://doi.org/10.22331/q-2021-05-14-456>
- [38] Henrik R. Larsson, Huanchen Zhai, Cyrus J. Umrigar, and Garnet Kin-Lic Chan. 2022. The chromium dimer: closing a chapter of quantum chemistry. <https://doi.org/10.48550/ARXIV.2206.10738>
- [39] Wim Lavrijsen, Ana Tudor, Juliane Müller, Costin Iancu, and Wibe de Jong. 2020. Classical Optimizers for Noisy Intermediate-Scale Quantum Devices. In *2020 IEEE International Conference on Quantum Computing and Engineering (QCE)*. 267–277. <https://doi.org/10.1109/QCE49297.2020.00041>

- [40] Lisha Li, Kevin Jamieson, Giulia DeSalvo, Afshin Rostamizadeh, and Ameet Talwalkar. 2017. Hyperband: A novel bandit-based approach to hyperparameter optimization. *The Journal of Machine Learning Research* 18, 1 (2017), 6765–6816.
- [41] Ying Li and Simon C. Benjamin. 2017. Efficient Variational Quantum Simulator Incorporating Active Error Minimization. *Phys. Rev. X* 7 (Jun 2017), 021050. Issue 2. <https://doi.org/10.1103/PhysRevX.7.021050>
- [42] Carlos Ortiz Marrero, Mária Kieferová, and Nathan Wiebe. 2021. Entanglement Induced Barren Plateaus. arXiv:2010.15968 [quant-ph]
- [43] Jarrod R. McClean, Sergio Boixo, Vadim N. Smelyanskiy, Ryan Babbush, and Hartmut Neven. 2018. Barren plateaus in quantum neural network training landscapes. *Nature Communications* 9, 1 (Nov 2018). <https://doi.org/10.1038/s41467-018-07090-4>
- [44] Jarrod R McClean, Jonathan Romero, Ryan Babbush, and Alán Aspuru-Guzik. 2016. The theory of variational hybrid quantum-classical algorithms. *New Journal of Physics* 18, 2 (2016), 023023.
- [45] Kosuke Mitarai, Yasunari Suzuki, Wataru Mizukami, Yuya O. Nakagawa, and Keisuke Fujii. 2020. Quadratic Clifford expansion for efficient benchmarking and initialization of variational quantum algorithms. arXiv:2011.09927 [quant-ph]
- [46] Masoud Mohseni, Peter Read, Hartmut Neven, Sergio Boixo, Vasil Denchev, Ryan Babbush, Austin Fowler, Vadim Smelyanskiy, and John Martinis. 2017. Commercialize Quantum Technologies in Five Years. *Nature* 543 (2017), 171–174. <http://www.nature.com/news/commercialize-quantum-technologies-in-five-years-1.21583>
- [47] Nikolaj Moll, Panagiotis Barkoutsos, Lev S Bishop, Jerry M Chow, Andrew Cross, Daniel J Egger, Stefan Filipp, Andreas Fuhrer, Jay M Gambetta, Marc Ganzhorn, et al. 2018. Quantum optimization using variational algorithms on near-term quantum devices. *Quantum Science and Technology* 3, 3 (2018), 030503.
- [48] Prakash Murali, Jonathan M Baker, Ali Javadi-Abhari, Frederic T Chong, and Margaret Martonosi. 2019. Noise-adaptive compiler mappings for noisy intermediate-scale quantum computers. In *Proceedings of the Twenty-Fourth International Conference on Architectural Support for Programming Languages and Operating Systems*. 1015–1029.
- [49] Prakash Murali, David C McKay, Margaret Martonosi, and Ali Javadi-Abhari. 2020. Software mitigation of crosstalk on noisy intermediate-scale quantum computers. In *Proceedings of the Twenty-Fifth International Conference on Architectural Support for Programming Languages and Operating Systems*. 1001–1016.
- [50] Luigi Nardi, David Koeplinger, and Kunle Olukotun. 2019. Practical Design Space Exploration. arXiv:1810.05236 [cs.LG]
- [51] Michael A Nielsen and Isaac Chuang. 2002. Quantum computation and quantum information.
- [52] Joe O’Gorman and Earl T. Campbell. 2017. Quantum computation with realistic magic-state factories. *Physical Review A* 95, 3 (Mar 2017). <https://doi.org/10.1103/physreva.95.032338>
- [53] Tirthak Patel and Devesh Tiwari. 2020. CLITE: Efficient and QoS-Aware Co-Location of Multiple Latency-Critical Jobs for Warehouse Scale Computers. In *2020 IEEE International Symposium on High Performance Computer Architecture (HPCA)*. 193–206. <https://doi.org/10.1109/HPCA47549.2020.00025>
- [54] Taylor L. Patti, Khadijeh Najafi, Xun Gao, and Susanne F. Yelin. 2021. Entanglement devised barren plateau mitigation. *Phys. Rev. Research* 3 (Jul 2021), 033090. Issue 3. <https://doi.org/10.1103/PhysRevResearch.3.033090>
- [55] Alberto Peruzzo, Jarrod McClean, Peter Shadbolt, Man-Hong Yung, Xiao-Qi Zhou, Peter J Love, Alán Aspuru-Guzik, and Jeremy L O’Brien. 2014. A variational eigenvalue solver on a photonic quantum processor. *Nature communications* 5 (2014), 4213.
- [56] Kirk A Peterson, David Feller, and David A Dixon. 2012. Chemical accuracy in ab initio thermochemistry and spectroscopy: current strategies and future challenges. *Theoretical Chemistry Accounts* 131, 1 (2012), 1–20.
- [57] John Preskill. 2018. Quantum Computing in the NISQ era and beyond. *Quantum* 2 (2018), 79.
- [58] Gokul Subramanian Ravi, Kaitlin N. Smith, Pranav Gokhale, Andrea Mari, Nathan Earnest, Ali Javadi-Abhari, and Frederic T. Chong. 2021. VAQEM: A Variational Approach to Quantum Error Mitigation. arXiv:2112.05821 [quant-ph]
- [59] Joschka Roffe. 2019. Quantum error correction: an introductory guide. *Contemporary Physics* 60, 3 (Jul 2019), 226–245. <https://doi.org/10.1080/00107514.2019.1667078>
- [60] Jonathan Romero, Ryan Babbush, Jarrod R. McClean, Cornelius Hempel, Peter Love, and Alán Aspuru-Guzik. 2018. Strategies for quantum computing molecular energies using the unitary coupled cluster ansatz. arXiv:1701.02691 [quant-ph]
- [61] Elliott Rosenberg, Paul Ginsparg, and Peter L. McMahon. 2021. Experimental error mitigation using linear rescaling for variational quantum eigensolving with up to 20 qubits. *Quantum Science and Technology* (Nov 2021). <https://doi.org/10.1088/2058-9565/ac3b37>
- [62] Rohan Basu Roy, Tirthak Patel, Vijay Gadepally, and Devesh Tiwari. 2021. *Bliss: Auto-Tuning Complex Applications Using a Pool of Diverse Lightweight Learning Models*. Association for Computing Machinery, New York, NY, USA, 1280–1295. <https://doi.org/10.1145/3453483.3454109>
- [63] Ilya G. Ryabinkin, Scott N. Genin, and Artur F. Izmaylov. 2018. Constrained variational quantum eigensolver: Quantum computer search engine in the Fock space. arXiv:1806.00461 [physics.chem-ph]
- [64] Frederic Sauvage, Sukin Sim, Alexander A. Kunitsa, William A. Simon, Marta Mauri, and Alejandro Perdomo-Ortiz. 2021. FLIP: A flexible initializer for arbitrarily-sized parametrized quantum circuits. arXiv:2103.08572 [quant-ph]
- [65] Peter W. Shor. 1997. Polynomial-Time Algorithms for Prime Factorization and Discrete Logarithms on a Quantum Computer. *SIAM J. Comput.* 26, 5 (Oct 1997), 1484–1509. <https://doi.org/10.1137/s0097539795293172>
- [66] Kaitlin N Smith, Gokul Subramanian Ravi, Prakash Murali, Jonathan M Baker, Nathan Earnest, Ali Javadi-Abhari, and Frederic T Chong. 2021. Error Mitigation in Quantum Computers through Instruction Scheduling. arXiv preprint arXiv:2105.01760 (2021).
- [67] Armands Strikis, Dayue Qin, Yanzhu Chen, Simon C. Benjamin, and Ying Li. 2021. Learning-based quantum error mitigation. arXiv:2005.07601 [quant-ph]
- [68] Qiming Sun, Timothy C. Berkelbach, Nick S. Blunt, George H. Booth, Sheng Guo, Zhendong Li, Junzi Liu, James McClain, Elvira R. Sayfutyarova, Sandeep Sharma, Sebastian Wouters, and Garnet Kin-Lic Chan. 2017. The Python-based Simulations of Chemistry Framework (PySCF). arXiv:1701.08223 [physics.chem-ph]
- [69] Ryui Takagi, Suguru Endo, Shintaro Minagawa, and Mile Gu. 2021. Fundamental limits of quantum error mitigation. arXiv:2109.04457 [quant-ph]
- [70] Wei Tang, Teague Tomesh, Martin Suchara, Jeffrey Larson, and Margaret Martonosi. 2021. CutQC: using small Quantum computers for large Quantum circuit evaluations. *Proceedings of the 26th ACM International Conference on Architectural Support for Programming Languages and Operating Systems* (Apr 2021). <https://doi.org/10.1145/3445814.3446758>
- [71] Swamit S Tannu and Moinuddin K Qureshi. 2019. Not all qubits are created equal: a case for variability-aware policies for NISQ-era quantum computers. In *Proceedings of the Twenty-Fourth International Conference on Architectural Support for Programming Languages and Operating Systems*. 987–999.
- [72] Kristian Temme, Sergey Bravyi, and Jay M Gambetta. 2017. Error mitigation for short-depth quantum circuits. *Physical review letters* 119, 18 (2017), 180509.
- [73] Jules Tilly, Hongxiang Chen, Shuxiang Cao, Dario Picozzi, Kanav Setia, Ying Li, Edward Grant, Leonard Wossnig, Ivan Rungger, George H. Booth, and Jonathan Tennyson. 2021. The Variational Quantum Eigensolver: a review of methods and best practices. arXiv:2111.05176 [quant-ph]
- [74] Justin M Turney, Andrew C Simmonett, Robert M Parrish, Edward G Hohenstein, Francesco A Evangelista, Justin T Fermann, Benjamin J Mintz, Lori A Burns, Jeremiah J Wilke, Micah L Abrams, et al. 2012. Psi4: an open-source ab initio electronic structure program. *Wiley Interdisciplinary Reviews: Computational Molecular Science* 2, 4 (2012), 556–565.
- [75] A V Uvarov and J D Biamonte. 2021. On barren plateaus and cost function locality in variational quantum algorithms. *Journal of Physics A: Mathematical and Theoretical* 54, 24 (May 2021). <https://doi.org/10.1088/1751-8121/abfac7>
- [76] Steven Vancollie, Per Åke Malmqvist, and Valera Veryazov. 2016. Potential energy surface of the chromium dimer re-revisited with multiconfigurational perturbation theory. *Journal of chemical theory and computation* 12, 4 (2016), 1647–1655.
- [77] Victor Veitch, S A Hamed Mousavian, Daniel Gottesman, and Joseph Emerson. 2014. The resource theory of stabilizer quantum computation. *New Journal of Physics* 16, 1 (Jan 2014), 013009. <https://doi.org/10.1088/1367-2630/16/1/013009>
- [78] Samson Wang, Piotr Czarnik, Andrew Arrasmith, M. Cerezo, Lukasz Cincio, and Patrick J. Coles. 2021. Can Error Mitigation Improve Trainability of Noisy Variational Quantum Algorithms? arXiv:2109.01051 [quant-ph]
- [79] Samson Wang, Enrico Fontana, Marco Cerezo, Kunal Sharma, Akira Sone, Lukasz Cincio, and Patrick J Coles. 2020. Noise-induced barren plateaus in variational quantum algorithms. arXiv preprint arXiv:2007.14384 (2020).
- [80] Yu Zhang, Lukasz Cincio, Christian FA Negre, Piotr Czarnik, Patrick Coles, Petr M Anisimov, Susan M Mniszewski, Sergei Tretiak, and Pavel A Dub. 2021. Variational quantum eigensolver with reduced circuit complexity. arXiv preprint arXiv:2106.07619 (2021).
- [81] Yiqing Zhou, E. Miles Stoudenmire, and Xavier Waintal. 2020. What Limits the Simulation of Quantum Computers? *Phys. Rev. X* 10 (Nov 2020), 041038. Issue 4. <https://doi.org/10.1103/PhysRevX.10.041038>