Improving the Variational Quantum Eigensolver Using Variational Adiabatic Quantum Computing

STUART M. HARWOOD, DIMITAR TRENEV, and SPENCER T. STOBBER, ExxonMobil
Corporate Strategic Research, USA
PANAGIOTIS BARKOUTSOS, IBM Quantum, IBM Research Zurich, Switzerland
TANVI P. GUJARATI, IBM Quantum, IBM Research Almaden, USA
SARAH MOSTAME and DONNY GREENBERG, IBM Quantum, IBM T. J. Watson Research
Center, USA

The variational quantum eigensolver (VQE) is a hybrid quantum-classical algorithm for finding the minimum eigenvalue of a Hamiltonian that involves the optimization of a parameterized quantum circuit. Since the resulting optimization problem is in general nonconvex, the method can converge to suboptimal parameter values that do not yield the minimum eigenvalue. In this work, we address this shortcoming by adopting the concept of variational adiabatic quantum computing (VAQC) as a procedure to improve VQE. In VAQC, the ground state of a continuously parameterized Hamiltonian is approximated via a parameterized quantum circuit. We discuss some basic theory of VAQC to motivate the development of a hybrid quantum-classical homotopy continuation method. The proposed method has parallels with a predictor-corrector method for numerical integration of differential equations. While there are theoretical limitations to the procedure, we see in practice that VAQC can successfully find good initial circuit parameters to initialize VQE. We demonstrate this with two examples from quantum chemistry. Through these examples, we provide empirical evidence that VAQC, combined with other techniques (an adaptive termination criteria for the classical optimizer and a variance-based resampling method for the expectation evaluation), can provide more accurate solutions than “plain” VQE, for the same amount of effort.

CCS Concepts: • **Mathematics of computing** → **Nonconvex optimization**; • **Computing methodologies** → *Quantum mechanic simulation*; • **Theory of computation** → **Quantum computation theory**

Additional Key Words and Phrases: Variational methods, adiabatic computing

ACM Reference format:

Stuart M. Harwood, Dimitar Trenev, Spencer T. Stober, Panagiotis Barkoutsos, Tanvi P. Gujarati, Sarah Mostame, and Donny Greenberg. 2022. Improving the Variational Quantum Eigensolver Using Variational Adiabatic Quantum Computing. *ACM Trans. Quantum Comput.* 3, 1, Article 1 (January 2022), 20 pages.
<https://doi.org/10.1145/3479197>

Authors' addresses: S. M. Harwood (corresponding author), D. Trenev, and S. T. Stober, ExxonMobil Corporate Strategic Research, 1545 US-22, Annandale, NJ, 08801; email: stuart.m.harwood@exxonmobil.com; P. Barkoutsos, IBM Quantum, IBM Research Zurich, Säumerstrasse 4, 8803 Rüschlikon, CH, Switzerland; T. P. Gujarati, IBM Quantum, IBM Research Almaden, 650 Harry Road, San Jose, CA, 95120; S. Mostame and D. Greenberg, IBM Quantum, IBM T. J. Watson Research Center, 1101 Kitchawan Road, Yorktown Heights, NY, 10598.

Permission to make digital or hard copies of part or all 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 third-party components of this work must be honored. For all other uses, contact the owner/author(s).

© 2022 Copyright held by the owner/author(s).

2643-6817/2022/01-ART1 \$15.00

<https://doi.org/10.1145/3479197>

1 INTRODUCTION

The **variational quantum eigensolver (VQE)** [27, 30, 34] is a hybrid quantum-classical algorithm for approximating the minimum eigenvalue of a given Hamiltonian. A parameterized quantum circuit, or *ansatz*, is tuned via a classical optimization method to minimize the expected value of the Hamiltonian. This expected value is estimated on a quantum computer, and thus this evaluation is noisy (at the very least, there is noise from the finite number of circuit samples or shots used to estimate the expected value). Consequently, VQE has characteristics of a stochastic optimization problem. VQE has been used to estimate the ground-state molecular energy of small molecules [20, 34], and as a heuristic for Ising problems [30, 33].

While VQE has shown promise in some small size problems, its ability to scale to larger problems is an open question; ultimately, the classical optimization methods on which VQE relies are at best local methods, while the objectives are typically nonconvex with multiple local minima. Convergence to a local, but not global, minimizer limits the accuracy of the estimated eigenvalue.

To address this fundamental shortcoming of VQE, as well as the parametric nature of certain applications, we adopt the concept of **variational adiabatic quantum computing (VAQC)**. In this approach, we use VQE to calculate the minimum eigenvalues of a parameterized Hamiltonian, taking advantage of the continuity of the problems to bootstrap or warm-start the local optimization methods. At a high level, we see that VAQC uses a parameterized Hamiltonian to navigate an energy landscape to arrive at the minimum eigenvalue of a target Hamiltonian. We draw connections to predictor-corrector methods for integrating differential equations and numerical continuation methods [3] by recasting the problem in a differential setting. We introduce a few other practical numerical techniques (adaptive termination of the classical optimization procedure and resampling) for improving the performance of VQE.

As the name suggests, we borrow notions and ideas from **adiabatic quantum computation (AQC)** and adiabatic simulation; see for instance Reference [2] for a review. Improvements to adiabatic simulation for molecular systems was explored recently in Reference [24]. However, the present work focuses on a computational protocol for the enhancement of variational algorithms used on near term gate-based quantum computers. As we discuss in Section 2, in AQC a quantum state evolves in some Hilbert space via the Schrödinger equation. In contrast, in VAQC a quantum state evolves in some reduced space (due to the *ansatz*) via a series of applications of VQE. A similar method was proposed in Reference [17], where the authors introduce the idea of adiabatically assisted VQE. The present work expands on this idea, provides different theoretical perspectives, and discusses more connections to other numerical methods.

An approach was recently introduced in Reference [11] under the name “adiabatic variational quantum computing.” This approach is essentially a variational time evolution method, a hybrid quantum-classical method for simulating the Schrödinger equation. However, our approach to VAQC in the present work is fundamentally different; we do not try to simulate unitary evolution. We provide a different theoretical perspective to motivate VAQC and to establish certain continuity properties of the optimal *ansatz* parameters. Further, this alternate point of view permits us to propose a numerical method that is different from typical numerical integration methods, and in particular the method proposed in Reference [11].

Compared to the quantum approximate optimization algorithm [16] and other work like [29], the VAQC approach described here permits in practice a wider range of *ansatz* forms. Inspired by the adiabatic theorem, the main requirement on the *ansatz* is that it can represent the ground state of the initial Hamiltonian (and, critically, the corresponding values of the *ansatz* parameters are known). If there is flexibility in the problem to do so, then these considerations suggest designing the *ansatz* and the initial Hamiltonian together to achieve this goal. In the examples that we

consider, we still use “hardware-efficient” ansätze [20], and suggest a simple modification to how these ansätze are typically constructed.

The VAQC approach and techniques for augmenting VQE are independent of the classical optimization method (within reason—ultimately we still need noise-robust optimization methods). Indeed, we demonstrate that these techniques can practically improve the performance of VQE (for instance, an order of magnitude or more reduction in error) using a few different optimization methods (namely stochastic gradient descent, simultaneous perturbation stochastic approximation, and the Nakanishi-Fujii-Todo method). While Reference [41] notes that the performance of these methods can be sensitive to the hyperparameters used (like step sizes), we expect any hyperparameter tuning would benefit the augmented procedures as well, and only further improve the results we obtain. Ultimately, the techniques we propose offer similarly or more accurate answers, for the same number of samples. When considering the number of unique circuits evaluated, these techniques can offer an order of magnitude reduction.

Meta-VQE, recently proposed in Reference [10], has a similar aim of efficiently calculating the ground-state energies of a parameterized Hamiltonian. However, the approach in meta-VQE is to prepend to the variational ansatz another parameterized circuit that aims to “learn” the parametric form of the Hamiltonian. In this work, we focus on a Hamiltonian depending on a single parameter (“time”) and are less concerned with fitting or learning a functional form of the ground-state energy.

As mentioned, VQE has been applied to Ising problems and similar classical combinatorial optimization problems. Although we do not consider such examples, the present work could also be applied to these problems. The recent work in Reference [15] has a similar goal of improving variational algorithms for combinatorial optimization through warm-start strategies. Combining those ideas with the present work is a subject for future research.

This work is organized as follows. In Section 2, we begin with a few theoretical considerations to describe and motivate VAQC. In Section 3, we describe its numerical implementation as a bootstrapping procedure, along with the other practical improvements like resampling and adaptive termination. In Section 4, we discuss some numerical results with simulated quantum devices. In particular, we show how VAQC can improve on the performance of plain VQE to find a better quality ground-state energy of a small molecule. Further, we demonstrate that in certain settings, like calculating the Born-Oppenheimer potential energy surface, VAQC and the proposed practical improvements can yield more consistently accurate results. We then conclude with a few final thoughts.

2 VARIATIONAL ADIABATIC QUANTUM COMPUTING

Abstractly, the goal is to calculate

$$f^*(t) = \min_{\theta} f(\theta, t) \equiv \langle \psi(\theta) | H(t) | \psi(\theta) \rangle \quad (1)$$

for a range of values of t . Here, H is a parameterized Hamiltonian while ψ is the VQE ansatz, a parameterized quantum state (trial wave function). For simplicity, we assume that t is a single real value (typically time), and for expositional purposes at least, we assume that it is scaled and shifted so $t \in [0, 1]$. To be concrete, we assume a system of n qubits, so we define the Hilbert space of interest as $\mathcal{H} \equiv (\mathbb{C}^2)^{\otimes n}$, and so H is a mapping from $[0, 1]$ to the space of self-adjoint linear operators on $\mathcal{H}$, while ψ is a mapping from some domain $\mathcal{D}$ to $\mathcal{H}$. We typically think of $\mathcal{D}$ as a subset of $\mathbb{R}^p$ (that is, the quantum state is parameterized by p real parameters); however, we will see that some flexibility is useful.

We may consider Problem (1) as an approximation of an adiabatic computation; in this setting $H(t) = (1-t)H_I + tH_T$, where H_I is an initial (or mixing) Hamiltonian, H_T is the target Hamiltonian,

and the goal is to determine the ground state of H_T . As another application of Equation (1), consider calculating the potential energy surface of a molecule, where H is the Hamiltonian parameterized by some molecular coordinate t .

At its core, VAQC may be seen as a hybrid quantum-classical homotopy/numerical continuation method, with the goal of approximating the ground state of the parametric Hamiltonian at each instant in time. At a high level, the same can be said of AQC. In VAQC, the approximation to the ground state is in a reduced space due to the ansatz, and is evolved by a hybrid quantum-classical method (e.g., VQE). In AQC, however, the approximation to the ground state occurs in the full Hilbert space and evolves via unitary evolution (the Schrödinger equation).

The basic requirement of VAQC is that the optimal ansatz parameters θ^* are continuous as a function of t . Consequently, the solution of Problem (1) for some value of t will be close to the solution of Problem (1) for some other nearby value t' . This proves very useful when solving Problem (1) by local optimization methods like gradient descent.

To establish the required continuity properties, we begin with the following result, stating that there exists a continuous ground state of H . See Appendix A for its proof. This result requires a continuously parameterized Hamiltonian with a strictly positive minimum gap between the ground-state energy and first excited-state energy. As another, more technical, connection with AQC, similar conditions appear in versions of the adiabatic theorem [19].

PROPOSITION 1. *Assume H is continuously differentiable on $[0, 1]$. Assume that there exists real constant $\delta > 0$ such that for each t , the lowest and second lowest eigenvalues (of $H(t)$), $\lambda_0(t)$ and $\lambda_1(t)$ respectively, are separated by δ : $\lambda_0(t) + \delta < \lambda_1(t)$. Then there exists a continuous mapping $\phi^* : [0, 1] \rightarrow \mathcal{H}$ such that $\phi^*(t)$ is a ground state of $H(t)$ for all $t \in [0, 1]$.*

The intuition is that if the ansatz ψ can represent the ground state $\phi^*(t)$ for each t , and this ground state is continuous, then the optimal ansatz parameters θ^* should be continuous as well. However, making this precise essentially requires the continuity of ψ^{-1} , the inverse mapping of ψ , which is not a property of ansätze that typically has been studied. Since we normally assume that ψ itself also is continuous, this leads to the assumption that ψ is a homeomorphism. Furthermore, the assumption that the ansatz can represent the ground state is nontrivial as well, but one that is common to most variational methods. Nevertheless, we state the following result to make it explicit what kinds of assumptions are needed.

PROPOSITION 2. *Assume that the ansatz $\psi : \mathcal{D} \rightarrow \mathcal{H}_\psi$ is a homeomorphism for some subset $\mathcal{H}_\psi \subset \mathcal{H}$. Assume that there exists a continuous mapping ϕ^* such that $\phi^*(t) \in \mathcal{H}_\psi$ and is a ground state of $H(t)$ for all $t \in [0, 1]$. Then there exists a continuous mapping $\theta^* : [0, 1] \rightarrow \mathcal{D}$ such that $\theta^*(t)$ is an optimal solution to Problem (1) for each $t \in [0, 1]$.*

PROOF. As a homeomorphism, ψ^{-1} is a continuous mapping $\mathcal{H}_\psi \rightarrow \mathcal{D}$. Thus define $\theta^*(t) = \psi^{-1}(\phi^*(t))$, and we see that it must be a solution of Problem (1) and θ^* is continuous. $\square$

Although we typically think of ψ as a mapping on $\mathbb{R}^p$, the abstraction of its domain to the space $\mathcal{D}$ is useful in the preceding result. For instance, consider an ansatz whose parameterization is via rotation gates and is periodic with period 2π with respect to each parameter. Consequently, as a function on $\mathbb{R}^p$, such an ansatz cannot be bijective (and thus it cannot be a homeomorphism). If we restrict the parameters to $[0, 2\pi]^p$, then we potentially impose discontinuities on the inverse map. However, if we define $\mathcal{D}$ as the quotient space of $\mathbb{R}^p$ with respect to the equivalence relation defined by this periodicity¹ (along with the corresponding quotient topology), then we have

¹Specifically, the equivalence relation in this example is $\theta \sim \theta'$ iff for each $i \in \{1, \dots, p\}$ there exists $k_i \in \mathbb{Z}$ such that $\theta_i + k_i 2\pi = \theta'_i$.

some hope that this ansatz could be a bijection, and subsequently a homeomorphism. Then the conclusion of Proposition 2 is that θ^* is continuous with respect to this quotient topology on $\mathcal{D}$; we can then define² a continuous mapping $\vartheta^* : [0, 1] \rightarrow \mathbb{R}^p$ such that $\vartheta^*(t)$ is a solution of Problem (1) for each t . Practically, this suggests in this setting that we should *not* impose bounds on the parameters when performing the optimization of Problem (1).

In AQC, the computation is initialized with a known ground state $\phi^*(0)$ of the initial Hamiltonian $H(0) = H_I$. Analogously, in VAQC, we require initial ansatz parameters θ^0 such that the ansatz represents the ground state of the initial Hamiltonian:

$$\psi(\theta^0) = \phi^*(0).$$

If ψ is a homeomorphism, then as in the proof of Proposition 2 we must have $\theta^0 = \psi^{-1}(\phi^*(0)) = \theta^*(0)$.

The requirement that the ansatz be able to represent the ground state at each instant in time may seem too strong, especially if the goal is to find the ground state of the final target Hamiltonian H_T only. Depending on the properties of the ansatz, a modified analysis is possible. For instance, if the ansatz spans a linear subspace (or, more accurately, the range $\psi(\mathcal{D})$ of the ansatz equals $\mathcal{L} \cap \mathcal{S}$, where $\mathcal{L}$ is some linear subspace of $\mathcal{H}$ and $\mathcal{S}$ is the unit sphere), then we are really interested in the spectral gap of the reduced Hamiltonian $A^\dagger H(t)A$, where A is some linear operator from a reduced dimension space $\mathbb{C}^d$ to $\mathcal{H}$ (specifically, its columns form an orthonormal basis for $\mathcal{L}$). Assuming that the reduced parameterized Hamiltonian $t \mapsto A^\dagger H(t)A$ satisfies conditions analogous to those in Proposition 1, then there exists a continuous ground state $\phi_{\mathcal{L}}^*$ of the reduced Hamiltonian. If $\mathcal{L}$ contains the ground state of H_T , then once again VAQC provides a method to find it.

Finally, since the motivation behind VAQC is to avoid sub-optimal local minimizers of Problem (1), one question is whether the optimal ansatz parameters θ^* are locally unique, and if so, how close do they come to another local minimizer. In the numerical implementation of VAQC, we must discretize the parameter t and take finite steps, and so the local uniqueness of the optimal ansatz parameters has implications for the size of the allowable time steps. While we leave a complete discussion for future work, we state a brief result that is relevant.

PROPOSITION 3. *Assume that we can identify $\mathcal{D}$ with some open subset of $\mathbb{R}^p$, f is twice-continuously differentiable, and $\theta^* : [0, 1] \rightarrow \mathcal{D}$ is continuous. Assume that for each $t \in [0, 1]$, $\theta^*(t)$ is a solution of Problem (1) and the Hessian matrix $\nabla_{\theta}^2 f(\theta^*(t), t)$ is positive definite. Then (for any norm $\|\cdot\|$) there exists a positive constant ϵ such that for all t , no other local minimizer of $f(\cdot, t)$ comes within a distance less than ϵ of $\theta^*(t)$.*

PROOF. Define $A(t) \equiv \nabla_{\theta}^2 f(\theta^*(t), t)$. Then A is positive definite-valued and continuous, and we can find a positive η such that for all t we have $\eta \leq (2\|A(t)^{-1}\|)^{-1}$. For any $\alpha > 0$, the set $\mathcal{K} = \{(\theta, t) : \|\theta - \theta^*(t)\| \leq \alpha, t \in [0, 1]\}$ is compact. Thus $\nabla_{\theta}^2 f$ is uniformly continuous on $\mathcal{K}$, and so we can find an $\epsilon > 0$ such that for all t and θ with $\|\theta - \theta^*(t)\| < \epsilon$ we have $\|\nabla_{\theta}^2 f(\theta, t) - A(t)\| < \eta$.

Now, for each t , $\nabla_{\theta} f(\theta^*(t), t) = 0$ by the necessary conditions for optimality [5, Prop. 1.1.1]. For each $t \in [0, 1]$ let $g_t : \theta \mapsto \theta - A(t)^{-1} \nabla_{\theta} f(\theta, t)$. We see that $\nabla_{\theta} f(\theta, t) = 0$ if and only if θ is a fixed point of g_t . Then for all t , $\nabla g_t(\theta) = (A(t) - \nabla_{\theta}^2 f(\theta, t))A(t)^{-1}$, and so $\|\nabla g_t(\theta)\| < \eta\|A(t)^{-1}\| \leq 1/2$ for all θ such that $\|\theta - \theta^*(t)\| < \epsilon$. It follows that g_t is a contraction mapping on $\mathcal{B}_{\epsilon}(\theta^*(t))$, the open ball of radius ϵ centered at $\theta^*(t)$. Thus g_t has at most one fixed point in $\mathcal{B}_{\epsilon}(\theta^*(t))$ [36, Thm. 9.23]; we conclude that $\theta^*(t)$ is the only value at which the gradient of f is zero. Thus, for all t , $\theta^*(t)$ is the only local minimizer, and further the only stationary point of $f(\cdot, t)$, in $\mathcal{B}_{\epsilon}(\theta^*(t))$. $\square$

²In this case, the quotient map induced by the equivalence relation is a covering map of $\mathcal{D}$, and so the path lifting property holds; we can “lift” θ^* to ϑ^* . See Sections 53 and 54 of Reference [31], in particular Lemma 54.1.

3 IMPLEMENTATION AND PRACTICAL IMPROVEMENTS

In this section, we describe the implementation of VAQC as bootstrapping and the techniques of adaptive termination and resampling in detail. While the previous section gives precise conditions under which VAQC has some theoretical guarantees, moving forward we will largely view VAQC from a more practical, heuristic perspective.

3.1 Bootstrapping

Again, recall that the motivation behind VAQC is to overcome the limitations of VQE when faced with a high-dimensional, nonconvex energy landscape with multiple local minima and stationary points. The proposed bootstrapping procedure aims to navigate this landscape through the solution of a series of related optimization problems, using the solution of the previous problem to warm-start or “bootstrap” the solution of the next problem. Specifically, we assume that we have an ordered sequence of values of t , $(t_1, t_2, \dots, t_K)$, at which f^* is wanted. Then for each k , if $\tilde{\theta}^k$ is an approximate solution of $\min_{\theta} f(\theta, t_k)$, we use $\tilde{\theta}^k$ as the initial point when solving $\min_{\theta} f(\theta, t_{k+1})$. Further, as discussed in Section 2, we assume that we have initial optimal parameters $\theta^0 \in \arg \min_{\theta} f(\theta, 0)$ that we use as the initial guess when solving the problem at t_1 . Based on Propositions 1 and 2, we hope that the optimal parameters are continuous, and that $\tilde{\theta}^k$ is a reasonable initial estimate for the next step.

To draw connections to numerical continuation methods [3], we discuss VAQC from the point of view of integrating differential equations. From this perspective, we transform the problem of finding a minimizer of Problem (1) for each t , into the problem of finding a solution of the necessary conditions for optimality: $\frac{\partial f}{\partial \theta_i}(\theta, t) = 0$ for each $i \in \{1, \dots, p\}$ and each $t \in [0, 1]$. We then differentiate these equations with respect to t and look for a solution θ^* of the implicit ordinary differential equations

$$\sum_{j=1}^p \frac{\partial^2 f}{\partial \theta_j \partial \theta_i}(\theta^*(t), t) \frac{\partial \theta_j^*}{\partial t}(t) + \frac{\partial^2 f}{\partial t \partial \theta_i}(\theta^*(t), t) = 0, \quad \forall (i, t) \in \{1, \dots, p\} \times [0, 1]. \quad (2)$$

The initial conditions for the differential equations are set to $\theta^*(0) = \theta^0$. In general (without the supporting theory of the previous section), a solution of (2) at best gives a *local* minimizer at any t . Similar equations are derived in Reference [28]; as in that work, we could use an explicit Euler method to numerically integrate Equations (2). Variational (real or imaginary) time evolution also relies on the numerical integration of implicit ordinary differential equations in the space of the ansatz parameters, and in practice often use an explicit Euler method [11, 25, 43].

The idea of integrating differential equations to find the zero of a system of algebraic equations is the basic idea in numerical continuation methods [3]. A common approach there, however, is to use predictor-corrector methods for numerical integration. The proposed bootstrapping method is akin to using $\tilde{\theta}^k$ as a zeroth-order predictor of the solution at t_{k+1} , followed by a local optimization method as the corrector iteration to improve this estimate.

Since the corrector iteration is in effect an application of VQE, it can be expensive, measured in terms of the number of objective function evaluations. However, an Euler step requires an accurate estimate of the Hessian matrix of f , and we will see in practice (see discussion in Section 4.2) that the cost of the corrector iteration is comparable to the cost of an explicit Euler step. Meanwhile, compared to an explicit Euler method, a benefit of the bootstrapping method is that fairly large time steps can be taken, since errors in the approximation tend to be corrected by the solution of the optimization problem, instead of accumulating as in Euler’s method. Further, depending on the application, the ansatz parameter estimates $\tilde{\theta}^k$ do not need to be highly accurate; we only need them to be in the “basin of attraction” of the global optimal values. For instance, in the setting of

an adiabatic computation, we only need the *final* ground-state energy with high accuracy, and the VAQC approach is a heuristic to provide us with a high-quality initial point $\tilde{\theta}^K$ that can be used in a final application of VQE to the target Hamiltonian H_T .

3.2 Adaptive Termination

Intuitively, if either the ansatz parameters or the objective estimates are varying below some tolerance, then we have determined the objective or the wave function as accurately as we need, and there is no need for further optimization. This is the motivation behind adaptive termination. This is a standard technique in deterministic optimization, although in the stochastic setting some modifications are necessary.

Since most optimization methods tend to incrementally improve a point, moving toward a minimum, the idea behind adaptive termination is to look at the change in a windowed average of past iterates. The candidate ansatz parameter values θ^k and objective estimates $\hat{f}^k$ from the last N_w iterations are saved (for some given value of N_w). The method terminates when we have either

$$\frac{1}{N_w} \left| \sum_{i=1}^{N_w} \hat{f}^{k-i} - \sum_{i=1}^{N_w} \hat{f}^{k+1-i} \right| \leq \varepsilon_f \quad \text{or} \quad \frac{1}{N_w} \left\| \sum_{i=1}^{N_w} \theta^{k-i} - \sum_{i=1}^{N_w} \theta^{k+1-i} \right\| \leq \varepsilon_\theta,$$

for given tolerances ε_f and ε_θ . Evidently the terms in these windowed averages cancel out and the expressions simplify. In convex stochastic optimization, averaging the iterates is discussed in Reference [35].

Appropriate values of N_w , ε_f , and ε_θ depend on the problem (for instance, the scaling of the objective f influences the choice of the tolerance ε_f). In general, smaller values of ε_f and ε_θ lead to more accurate results, at the cost of more iterations of the optimization method. Meanwhile, larger values of N_w can help average out noise in the estimates.

3.3 Resampling

Again, the objective function of the optimization problem (1) is evaluated on a quantum computer, and thus we only have access to a noisy estimator. For simplicity, we assume this estimator $\hat{f}_m(\theta, t)$ is a sample average estimator consisting of m independent “samples.” The number of circuit evaluations and measurements that are actually required for each sample depends on the quantum hardware and the decomposition of $H(t)$ into basis gates (e.g., as a sum of tensor products of Pauli operators). We will largely ignore this, and when we refer to a “circuit evaluation,” “shot,” or “sample,” we are referring to some unit of quantum computational resource, m of which is required to calculate $\hat{f}_m(\theta, t)$. This ignores, for instance, that H may be more efficient to evaluate (e.g., have a simpler decomposition as Pauli operators) for some values of t than others. However, we feel this abstraction is acceptable, as it makes the present analysis independent of the specific estimation scheme used.

After the optimization/corrector iteration finishes at a parameter value t , we have (an approximation of) optimal ansatz parameters $\tilde{\theta}$. We still have the challenge of accurately estimating the value of $f(\tilde{\theta}, t)$; the base estimator $\hat{f}_m$ may not be sufficiently accurate.

Fortunately, in many cases we have access to an estimate of the variance of $\hat{f}_m(\tilde{\theta}, t)$. See for instance Section V.A of Kandala et al. [20] for the details of the variance estimator in Qiskit [1]. Assuming this variance is $\sigma^2(\tilde{\theta}, t)$, the average of N_r repeated independent estimators has variance $\sigma^2(\tilde{\theta}, t)/N_r$. Consequently, we can get an estimate of the objective with standard deviation less than or equal to ε_r by taking $N_r > \sigma^2(\tilde{\theta}, t)/\varepsilon_r^2$. In fact, considering that $\hat{f}_m$ already consists of m samples, we take $N_r > m\sigma^2(\tilde{\theta}, t)/\varepsilon_r^2$ and construct an estimator $\hat{f}_{N_r}(\tilde{\theta}, t)$ with this many samples in the first

place. See also Reference [27], which discusses similar ideas as well as an interesting approach from a Bayesian perspective.

However, we reiterate that this proposed resampling procedure is only applied to the final, optimized ansatz parameters. Another benefit of this construction of the estimator is that it adapts; based on the variance (which depends on t and θ), we use just enough samples to yield a final objective estimate within a desired tolerance. Using a noise-robust stochastic optimization method, few samples are needed during optimization when evaluating $\hat{f}_m(\theta, t)$. But when accuracy matters, at the final evaluation of the optimal objective, we use a high number of samples automatically determined by this resampling procedure.

4 EXPERIMENTS

In this section, we perform simulations to assess the effectiveness of VAQC and the improvements to VQE. These are performed with the QasmSimulator in Qiskit v0.15 [1]. In all experiments, we use $m = 64$ samples to build the objective estimator $\hat{f}_m(\theta, t)$. This is a relatively low number of samples and results in a fair amount of sampling noise. While not one of our main conclusions, this permits us to assess how noise-robust the various methods are.

The two examples we consider demonstrate two common applications of VQE in quantum chemistry: calculating a ground-state energy for a single molecular configuration and calculating an energy surface as a function of molecular coordinates. In the first case, we show how VAQC, with an artificially parameterized Hamiltonian, can yield more accurate energy values than plain VQE. In the second case, we take advantage of the naturally parametric Hamiltonian and apply VAQC to robustly calculate an energy surface.

4.1 Optimization Methods

We will use three optimization methods in our experiments: **stochastic gradient descent (SGD)**, **simultaneous perturbation stochastic approximation (SPSA)**, and **Nakanishi-Fujii-Todo (NFT)**. See Appendix B for more details. Each method is well suited to stochastic settings. Each method may be augmented with adaptive termination, bootstrapping, resampling, or a combination thereof; these are denoted with the suffixes “A,” “B,” or “R,” respectively. For instance, “NFT-BR” refers to NFT with bootstrapping and resampling applied.

There are many other optimization methods that can be applied in this context; see References [22, 41] for more extensive discussion along these lines. Again, our intention is not to exhaustively compare or benchmark optimization methods, but rather to assess the effect of VAQC and the other augmentations. The three methods are meant to be representative of a few different classes of optimization methods: **strongly local, gradient-based methods (SGD)**; derivative-free methods with some random search (SPSA); and domain-specific methods with some guarantee of finding a global minimum (NFT).

4.2 VAQC Calculation of Ground State of Lithium Hydride

We use the VAQC approach to robustly compute the ground-state energy, and compare this with VQE. The target Hamiltonian H_T acting on four qubits comes from the electronic energy of LiH when the distance between the nuclei is 2.5 Å. This is not the equilibrium bond length; we specifically choose this internuclear distance, since it can be a challenging point for VQE to find a ground-state energy. See Appendix C for its full specification. Errors are with respect to the full configuration interaction energy (i.e., the lowest eigenvalue of the Hamiltonian, calculated with classical methods). The objective function is measured in units of Hartree, and so we choose the tolerances with the aim of achieving an accuracy of order 10^{-3} . Thus, for instance, we will use $\epsilon_r = 5 \times 10^{-4}$ in the resampling procedure, whenever it is used.

Fig. 1. Ansatz for LiH ground-state study.

Procedure. We use a four-layer RY ansatz with a few additional gates appended to the end of the circuit. This results in an ansatz with $p = 20$ parameters (see Figure 1). Note that $\psi(0) = |0110\rangle$, since $R_Y(0) = I$, each CNOT does not modify the initial state $|0000\rangle$, and the final X gates flip the second and third qubits. We define the initial Hamiltonian as

$$H_I = (0.568)I^{\otimes 4} + (-0.102)Z_1 + (0.245)Z_2 + (0.102)Z_3 + (-0.245)Z_4,$$

which indeed has $\psi(0) = |0110\rangle$ as a ground state. The motivation for this form of H_I is that the target Hamiltonian as a weighted sum of tensor products of Pauli operators has the form $H_T = H_I + (\text{other terms})$, where “other terms” captures the rest of the terms that do not include Z_1, Z_2, Z_3, Z_4 , or $I^{\otimes 4}$. Furthermore, all other terms have a weight with absolute value less than 0.2, and so H_I is capturing some of the more significant terms. While the Hartree-Fock state and corresponding Hamiltonian are another option for constructing the ansatz and initial Hamiltonian, the present choice demonstrates that even a mildly problem-specific heuristic for constructing the initial Hamiltonian is successful.

Subsequently, we define the parameterized Hamiltonian: $H(t) = (1 - t)H_I + tH_T$. The “time” points are taken to be $t_k = 1 - (1 - 0.05k)^3$ for $k \in \{1, \dots, 20\}$; this is akin to using a dimensionless time step of 0.05, along with a schedule of $s : t \mapsto 1 - (1 - t)^3$. While the term “schedule” is borrowed from adiabatic computation, the motivation behind its form in the present setting is that we would like to approach the target Hamiltonian gradually, without taking any large steps that might inhibit the convergence of the optimization-based corrector iteration. Consequently, for a fixed number of time steps, we posit that they should be distributed so as to avoid large perturbations near the target Hamiltonian (this motivates a schedule of the form $t \mapsto 1 - (1 - t)^a$ with $a > 1$). This should also help with the relatively slow convergence of SGD, which we use for the corrector iteration.

The corrector iteration in specific is SGD-AB, with a maximum of $2p = 40$ iterations, $N_w = 10$, $\varepsilon_f = 10^{-3}$, and $\varepsilon_\theta = 10^{-3}$. SGD requires $2p + 1 = 41$ objective estimates per iteration (using the parameter shift rule to estimate the gradient). Thus, each corrector iteration requires up to $2p(2p + 1)$ objective estimates. This is roughly the same amount of effort required to estimate the Hessian matrix of f , which would be required for an Euler method applied to the differential equation (2): each of the $p(p + 1)/2$ unique entries of the symmetric Hessian require up to four objective estimates via an iterated parameter-shift rule to get a second derivative (although see Reference [28] for an alternative approach that requires fewer estimates if an ancilla qubit is available). However, in our approach, the adaptive termination means that we typically use fewer than the 40 allowed iterations.

Table 1. Accuracy of VAQC and VQE for LiH Ground-state Energy

Method	Mean absolute error (Ha)	Min / Max absolute error (Ha)
VAQC	0.0016	0.0004 / 0.0025
VQE	0.0640	0.0095 / 0.1641

With these settings, VAQC gets close to a good solution, but requires some polishing; consequently, at the last time point $t_{20} = 1$, we apply SGD (with adaptive termination, bootstrapping, and resampling) with a maximum of 80 iterations, $N_w = 10$, $\epsilon_f = 10^{-4}$, and $\epsilon_\theta = 5 \times 10^{-4}$.

VAQC will require more computation than “plain” VQE applied directly to the target Hamiltonian H_T . Given a budget of iterations, the best we can do with VQE is to repeatedly apply it with different randomly chosen initial points. Thus, we compare against plain VQE, using SGD-AR and the same ansatz, but with initial ansatz parameters uniformly randomly chosen from $[-\pi, \pi]^p$. We take $N_w = 10$, $\epsilon_f = 10^{-4}$, $\epsilon_\theta = 5 \times 10^{-4}$ (the same values used for the final time point of VAQC).

Results. VAQC robustly produces a high-quality solution; over five independent trials, VAQC yields a final objective value within 2.5×10^{-3} Hartree of the exact energy. See Table 1. Over these five repetitions, VAQC requires 1913 total iterations of SGD. Giving this budget of iterations to plain VQE, we are able to sample 33 different initial ansatz parameters, with an average of 58 iterations required to converge VQE. However, even these multiple applications of VQE with different initial ansatz parameters do not produce as accurate an answer; the minimum absolute error that is achieved is 9.5×10^{-3} Hartree. Furthermore, the results are not nearly as consistent, with errors as large as 0.1641 Hartree. This supports the conclusion that the VAQC procedure is capable of navigating a non-trivial landscape to improve upon the performance of plain VQE.

VAQC requires on average 19.13 iterations per time step, or $19.13 \times 41 = 784.33$ function evaluations per time step. As noted, an explicit Euler step would require at least $4^{(p+1)/2} = 840$ function evaluations per time step (if an ancilla qubit is not available). Thus, we see that VAQC (with SGD, at least) has a per time step cost comparable to an explicit Euler scheme.

Optimization landscape. We use this example as an opportunity to assess the optimization landscape and, in particular, the possibility of barren plateaus. As discussed in References [9, 26], this is the situation where there is a broad region of ansatz parameter values that are suboptimal and also lead to small gradients of the objective f . Noisy gradient estimates mean that there is no clear direction in which to update the parameters. As a result, a stochastic gradient method initialized on such a plateau resembles a random walk. Even worse, under fairly broad conditions on the ansatz and Hamiltonian, the probability of being in this region can be high, potentially approaching 1 exponentially quickly in the number of qubits. Thus, the probability of leaving the barren plateau via a random walk is exponentially low. This implies that randomly initializing the ansatz parameters leads to a low probability of even finding a local minimum. Further, it is not unreasonable that the present example could suffer from barren plateaus; our target Hamiltonian is a “global” cost function as defined in Reference [9], since its terms involve tensor products of (non-identity) Pauli operators on all qubits (see Table 3). In this situation, barren plateaus can be a problem for a general class of ansätze at very low depths.

To assess this, we plot the distribution of the norm of the objective function gradient at the initial parameter values used in VQE and VAQC; see Figure 2. In the case of VQE, these are the gradients of $f(\cdot, 1)$ evaluated at the 33 randomly chosen initial parameters, while for VAQC, these are the gradients evaluated at the initial parameters of each bootstrapping step, $\nabla f(\hat{\theta}^{k-1}, t_k)$ (recall the notation in Section 3.1). We see that for VQE, the random initial points lead to a broad distribution of gradient norms. In contrast, for VAQC, the gradients are smaller and more tightly

Fig. 2. Histogram of initial gradient norms for VAQC and VQE. The initial gradients of VQE are evaluated at the 33 uniformly randomly chosen initial points. The initial gradients of VAQC are for each of the bootstrapping steps over the five independent trials.

clustered, supporting our assertion that the sequence of solutions remain close to a minimum as the Hamiltonian evolves.

The fact that the initial gradients in VAQC are smaller and yet it still produces a more accurate answer than randomly initialized VQE suggests that a barren plateau is not to blame for the poor performance of VQE; the larger gradients at the random initial parameters indicate that the initial behavior of VQE is not like a random walk. While the random initial parameters lead to some type of suboptimal stationary point, the evidence does not support that this point is on an “exponentially large” plateau, otherwise we would expect the random initial gradients to have a tighter distribution of smaller norms.

Still, it is shown in Reference [9] that a “local” cost function, like the one resulting from the initial Hamiltonian H_I in this example, can avoid barren plateaus. It is an intriguing possibility that VAQC (with an appropriate ansatz) could avoid barren plateaus by interpolating between a local and global cost function. Further work is merited.

4.3 Potential Energy Surface of Molecular Hydrogen

In this problem, our goal is to calculate an accurate Born-Oppenheimer potential energy surface for H_2 . We will use this as an opportunity to assess the impact of the various proposals (bootstrapping/VAQC, adaptive termination, and resampling).

The parameter t corresponds to internuclear distance, and we seek the electronic energy at the points $\{0.4, 0.55, 0.7, 0.85, 1.2, 1.6, 1.9, 2.2, 2.6, 3.0, 4.0, 5.0\}$ (Å). We construct the Hamiltonian using PySCF [40] and the STO-3G basis set [12, 18], following the procedure outlined in Reference [20]. This Hamiltonian is then transformed into the particle/hole representation [4] and mapped to a qubit Hamiltonian using parity mapping [8, 38], which results in a two qubit representation. We use a one-layer RY ansatz with additional gates prepended that prepare the Hartree-Fock state.³ See Figure 3. Again, errors are with respect to the full configuration interaction energy (calculated classically). We target an accuracy level of 10^{-3} Hartree. Thus, for instance, we will use $\varepsilon_r = 5 \times 10^{-4}$ in the resampling procedure, whenever it is used. While in general, lower error is better, we are more interested in consistently low error, and further, a method that achieves error much lower than 10^{-3} could be seen as wasting effort.

³As discussed in the lithium hydride study (Section 4.2), arguably a better ansatz would be one with the gates that prepare the Hartree-Fock state appended to the *end* of the circuit. In fact, the methods considered do perform better with this ansatz, but it subsequently makes it harder to see the effects of the proposed improvements like bootstrapping.

Fig. 3. Ansatz for H_2 potential energy surface study.

Procedure. For any method using adaptive termination, we take $N_w = 10$, $\varepsilon_f = 10^{-4}$, and $\varepsilon_\theta = 5 \times 10^{-4}$. Any method using bootstrapping begins at the smallest distance and proceeds in order of increasing distance. For any method or iteration that is not bootstrapped, the initial guess for the ansatz parameters is zero. While $\psi(0)$ is not exactly the ground state for the initial Hamiltonian at $0.4 \text{ \AA}$, it is a reasonable starting point.

For each optimization method, we at least look at the performance of the baseline optimizer, optimizer plus resampling, and optimizer plus resampling plus bootstrapping. In the particular case of SGD, we also look at SGD plus resampling, bootstrapping, and adaptive termination (for NFT and SPSA, we use the current implementations in Qiskit, which would have to be modified to allow adaptive termination).

To assess the robustness of the methods, we perform five independent repetitions of the potential energy surface calculation for each method. To be as fair as possible in our comparisons of the various methods, we make sure that each method uses approximately the same number of total samples in the calculation of the potential energy surface, over all the independent repetitions. Since it is only SGD-ABR with its adaptive termination for which this budget of samples is not known beforehand, we run this method first and allocate its effort to the others. While we do not directly control how many samples are used for resampling, two methods that use resampling will use approximately the same total number of samples during the resampling phase.

Thus, we run SGD-ABR first, calculate the total number of samples used, and set an iteration limit (which is the same for each internuclear distance) for SGD, SPSA, and NFT so that they approximately use the same total number of samples (taking into account the different number of function calls per iteration for each method). Meanwhile, we take the number of samples used only during the optimization phase from SGD-ABR, and again set an iteration limit for SGD-R, SGD-BR, SPSA-R, SPSA-BR, NFT-R, and NFT-BR, so that the optimization phases of each method use the same number of samples (while the number of samples used in the resampling phases are roughly similar).

Improvements in accuracy. Figure 4 visualizes the errors, over all independent repetitions and internuclear distances, for each baseline optimization method and its augmented versions. Meanwhile, Figures 5 and 6 show the progression of mean absolute error at each internuclear distance for each method. In general, where there is error, we must distinguish between sampling error and failure of the optimization method to find optimal ansatz parameters. The resampled versions of the optimization methods help make this distinction. For instance, for SPSA and SGD, the addition of resampling does not make significant improvements (besides eliminating negative errors, which must be due to sampling error). This indicates that most of the error for those methods is due to poor quality ansatz parameters. Meanwhile, SPSA-BR improves a lot over SPSA and SPSA-R, and in particular is nearly an order of magnitude more accurate in the region of 1 to 2 $\text{\AA}$, indicating that bootstrapping yields a significant improvement. Similarly, we see that SGD improves significantly with bootstrapping and resampling. Looking at Figure 6, SGD-BR and SGD-ABR consistently achieve the desired threshold of 10^{-3} Hartree across the internuclear distances. Furthermore, SGD-BR and SGD-ABR are not limited to an accuracy of 10^{-3} Hartree; with a tighter tolerance ε_r and more optimization iterations we expect them to improve.

Fig. 4. Box and whisker plot for error (over independent repetitions and internuclear distances) for various methods for H_2 potential energy surface. Boxes represent middle quartiles, whiskers represent minimum to maximum, and the central line is the mean. Note the change in scaling at 0.05 Hartree on the y -axis.

Fig. 5. Mean absolute error (over five independent repetitions for each internuclear distance) versus distance. We have SPSA (left) and NFT (right) along with the corresponding augmented versions.

While mean absolute error appears similar between SGD-BR and SGD-ABR, *maximum* absolute error over all the distances and independent repetitions is 0.0128 Hartree for SGD-BR, and 0.0043 Hartree for SGD-ABR. See also Table 2. This is where adaptive termination proves valuable; SGD-ABR expends more effort when necessary to get closer to a desired accuracy level.

Figures 4 and 5 show that NFT is already a fairly robust optimization method, with resampling appearing to make most of the improvement over the base optimizer. This makes sense, since NFT can take large steps when updating the ansatz parameters. Thus, the final parameters depend less on the initial values given to NFT, making bootstrapping less effective.

Reduction in unique circuits. While we have kept the total number of samples the same in our comparisons, the number of *unique* circuits that are executed is an interesting metric to consider. In many hardware architectures, the overhead of preparing a circuit can be significant, and so once prepared, executing the same circuit many times is relatively efficient. So, consider Table 2. We see that, compared to the base optimizers, adding bootstrapping and resampling results in better mean

Fig. 6. Mean absolute error (over five independent repetitions for each internuclear distance) versus distance for SGD and its augmented versions.

Table 2. Accuracy and Cost for Various Methods for H_2 Potential Energy Surface

Method	Mean absolute error (Ha)	Max absolute error (Ha)	Mean unique circuit evaluations
SPSA	0.0174	0.0721	2341
SPSA-BR	0.0056	0.0337	205
NFT	0.0158	0.1180	2340
NFT-BR	0.0059	0.0397	204
SGD	0.0320	0.2040	2340
SGD-BR	0.0011	0.0128	198
SGD-ABR	0.0007	0.0043	203

Unique circuit evaluations are estimated from the number of function evaluations during optimization.

and maximum absolute error, as well as an order of magnitude fewer unique circuit evaluations. This is because while resampling contributes many objective evaluations, they are all at the same ansatz parameter values.

5 CONCLUSIONS

We have presented the concept of VAQC, its practical implementation as a bootstrapping procedure, and some other practical numerical techniques to improve VQE. When applied to the calculation of the ground-state energies of simple molecules, these techniques yield consistently more accurate results compared to a more basic implementation of VQE. This is particularly evident when calculating a potential energy surface, where a series of high accuracy energies are needed.

We have established some basic theoretical guarantees for the performance of VAQC. While the conditions under which these guarantees hold are difficult to verify, VAQC performs quite well as a heuristic for improving VQE. We expect the VAQC approach to benefit from further research in AQC; analysis of the spectral gaps of Hamiltonians of interest should impact the approach

discussed here, and homotopy maps could inspire novel ansätze. Meanwhile, VAQC may be a useful tool for better understanding or simulating adiabatic computations.

APPENDICES

A PROOF OF PROPOSITION 1

The following proof essentially elaborates on Remarks 5.10 and 4.3 in Reference [21, Chapter II]. A fundamental tool in that analysis is the resolvent operator

$$R(z, t) = (H(t) - zI)^{-1},$$

where z is a complex number and I is the identity on $\mathcal{H}$. Evidently the resolvent is only defined for (z, t) such that z is not an eigenvalue of $H(t)$. So, let $\mathcal{D}_R$ be the preimage under $(z, t) \mapsto H(t) - zI$ of the set of invertible operators. Then, as the preimage under a continuous mapping of an open set, $\mathcal{D}_R$ is an open subset of $\mathbb{C} \times [0, 1]$. Further, R is well-defined on $\mathcal{D}_R$. Since H and the inversion of an operator are continuous and further, differentiable, R is continuous and differentiable on $\mathcal{D}_R$ with derivative $\frac{\partial R}{\partial t}(z, t) = -R(z, t)\dot{H}(t)R(z, t)$. This shows that $\frac{\partial R}{\partial t}$ is continuous on $\mathcal{D}_R$. Importantly, for any s , $\frac{R(z, t) - R(z, s)}{t - s} \rightarrow \frac{\partial R}{\partial t}(z, s)$ pointwise in z as $t \rightarrow s$, and the convergence is uniform on any compact subset of $\mathbb{C}$ not containing an eigenvalue of $H(s)$.⁴ That $\mathcal{D}_R$ is open is critical here; if K is a compact subset of the complex plane that does not contain an eigenvalue of $H(s)$, then $K \times \{s\}$ is a subset of $\mathcal{D}_R$, and so $K \times \{t\}$, for t near s , is also a subset of $\mathcal{D}_R$.

With this, we can define an eigenprojection P as

$$P = -\frac{1}{2\pi i} \oint_{\Gamma} R(z, s) dz.$$

Here, P is the eigenprojection corresponding to all eigenvalues of $H(s)$ enclosed by the positively-oriented closed curve Γ in the complex plane. This characterization of the eigenprojection also appears in proofs of the adiabatic theorem; see for instance Reference [19]. Another basic fact we will use is that the eigenvalues of H may be chosen as continuous functions on $[0, 1]$, since H is continuous and the eigenvalues are real. This means that for t close to s , Γ encloses the “same” eigenvalues.

PROPOSITION. *Assume H is continuously differentiable on $[0, 1]$. Assume that there exists real constant $\delta > 0$ such that for each t , the lowest and second lowest eigenvalues (of $H(t)$), $\lambda_0(t)$ and $\lambda_1(t)$ respectively, are separated by δ : $\lambda_0(t) + \delta < \lambda_1(t)$. Then there exists a continuous mapping $\phi^* : [0, 1] \rightarrow \mathcal{H}$ such that $\phi^*(t)$ is a ground state of $H(t)$ for all $t \in [0, 1]$.*

PROOF. To start, we show that we have a continuously differentiable eigenprojection. Fix $s \in [0, 1]$. We may choose a closed curve Γ enclosing the lowest eigenvalue $\lambda_0(s)$, and not enclosing or passing through any other eigenvalue (this is possible by the assumption that the lowest eigenvalue is separated by δ). Define

$$P_0(s) = -\frac{1}{2\pi i} \oint_{\Gamma} R(z, s) dz.$$

Consequently, $P_0(s)$ is the eigenprojection corresponding to $\lambda_0(s)$. Critically, this expression continues to hold for all t in a neighborhood of s by the continuity of λ_0 . Then note that

$$\frac{P_0(t) - P_0(s)}{t - s} = -\frac{1}{2\pi i} \oint_{\Gamma} \frac{R(z, t) - R(z, s)}{t - s} dz$$

⁴A little more detail here: if R_{δ} is defined on $\mathcal{D}_R$ as the difference quotient when $t \neq s$, and as $\frac{\partial R}{\partial t}(z, s)$ for $t = s$, then it is a continuous function on $\mathcal{D}_R$. Thus $R_{\delta}(\cdot, t)$ must converge uniformly to $R_{\delta}(\cdot, s) = \frac{\partial R}{\partial t}(\cdot, s)$ on compact subsets K of $\mathbb{C}$ such that $K \times (s - \epsilon, s + \epsilon) \subset \mathcal{D}_R$, for some positive real ϵ .

for t sufficiently close to s . Taking the limit as t approaches s , we obtain $\dot{P}_0(s) = -1/2\pi i \oint_{\Gamma} \frac{\partial R}{\partial s}(z, s) dz$ by the uniform convergence of $\frac{R(z, t) - R(z, s)}{t - s}$ mentioned above. Once again, this expression continues to hold for t near s , so P_0 is continuously differentiable on a neighborhood of s , and since s was arbitrary, it is continuously differentiable on $[0, 1]$.

Next, we show that there exists a continuous U such that for all t , $U(t)$ is invertible and $P_0(t) = U(t)P_0(0)U^{-1}(t)$. Since P_0 is a projection, $(P_0(t))^2 = P_0(t)$ for all t , so differentiating yields

$$P_0\dot{P}_0 + \dot{P}_0P_0 = \dot{P}_0; \quad (3)$$

multiplying from the right and left by P_0 gives

$$P_0\dot{P}_0P_0 = 0. \quad (4)$$

Then let $Q(t) = \dot{P}_0(t)P_0(t) - P_0(t)\dot{P}_0(t)$ be the commutator of $\dot{P}_0$ and P_0 . We have that

$$\dot{P}_0 = QP_0 - P_0Q \quad (5)$$

using the definition of Q , Equation (3), and Equation (4). We construct U as the solution of the initial value problem in linear ordinary differential equations

$$\dot{V}(t) = Q(t)V(t), \forall t \in [0, 1], \quad V(0) = V_0. \quad (6)$$

In particular, for the initial value $V_0 = I$, we define the solution as U . We note that a general solution of Equation (6) may be given by $V(t) = U(t)V_0$. The inverse of U is the solution of the initial value problem

$$\dot{W}(t) = -W(t)Q(t), \forall t \in [0, 1], \quad W(0) = I.$$

For a solution W of the above, we see that $\frac{d(WU)}{dt} = W\dot{U} + \dot{W}U = WQU - WQU = 0$; thus $W(t)U(t)$ is constant and must equal its initial value $W(0)U(0) = I$ for all t . This shows that U is invertible. Then note that P_0U is a solution of Equation (6) with initial value $P_0(0)$:

$$\frac{d(P_0U)}{dt} = \dot{P}_0U + P_0\dot{U} = (\dot{P}_0 + P_0Q)U = Q(P_0U),$$

where we have used Equation (5) in the last equality. Since the solution of Equation (6) is unique for a given initial value, $P_0(t)U(t)$ must coincide with the general solution noted above $U(t)P_0(0)$. In other words $P_0(t)U(t) = U(t)P_0(0)$ for all t ; thus, we have $P_0(t) = U(t)P_0(0)U^{-1}(t)$ as desired.

Finally, let ϕ^0 be any (non-zero) vector in the eigenspace of $H(0)$ corresponding to $\lambda_0(0)$ (so $P_0(0)\phi^0 = \phi^0$). Then define $\phi(t) = U(t)\phi^0$. Note that

$$\begin{aligned} P_0(t)\phi(t) &= (U(t)P_0(0)U^{-1}(t))(U(t)\phi^0) \\ &= U(t)P_0(0)\phi^0 \\ &= U(t)\phi^0 = \phi(t). \end{aligned}$$

So we see that $\phi(t)$ is an eigenvector corresponding to the lowest eigenvalue of $H(t)$, for all $t \in [0, 1]$. Since $U(t)$ is invertible, $\phi(t)$ is non-zero, so we can normalize it to obtain $\phi^*(t) \equiv \frac{\phi(t)}{\|\phi(t)\|}$. We have that $\phi^*(t)$ is a well-defined, normalized ground state of $H(t)$ for each t , and ϕ^* is continuous. $\square$

B OPTIMIZATION METHODS

We describe in detail the optimization methods used in our experiments.

SGD is an optimization method defined by the iteration

$$\theta^{k+1} = \theta^k - \gamma^k g^k.$$

Here, γ^k is the step length, and g^k is the current gradient estimate. Gradients are obtained using the parameter shift rule [13, 37] (we use RY-type ansätze for ψ , for which the parameter shift rule is applicable). Along with evaluation of the objective estimate at the current parameter point, this makes the cost per iteration of the SGD method $2p + 1$ objective evaluations. Step lengths γ^k decrease to zero with an asymptotic behavior like $1/k$; this is consistent with most convergence theories [6]. We use a step length equal to 1 for the first ten iterations, then decrease it by setting it to $1/(k - 10)$, for $k > 10$.

The addition of a momentum term is a common variant of SGD. The idea is that this can help escape local minima; see References [14, 23] for various perspectives and references. Again, our point here is to show improvement in nearly any optimizer; we use a basic form of SGD for simplicity.

SPSA is a gradient-free stochastic optimization method [39]. It has been used successfully with VQE in previous studies [20]. Each iteration consists of two objective evaluations, essentially giving a derivative approximation in a particular coordinate direction. The algorithm randomizes these search directions and carefully controls the sequence of steplengths that it takes. Hyperparameter settings are the default in Qiskit, with the exception of the initial step size that is set to 1.0.

NFT is a sequential coordinate minimization method [32]. When ψ is an RY-type ansatz, the objective f is sinusoidal as a function of a single parameter θ_i with the others fixed. The method leverages this fact to estimate f as a function of a single parameter and analytically minimize it in this coordinate direction. NFT requires two to three objective evaluations per iteration (the frequency of the extra evaluation is a hyperparameter—we use the default in its implementation in Qiskit). In the noise-free setting, at least, sequential minimization methods enjoy reasonable convergence guarantees [42].

C LITHIUM HYDRIDE HAMILTONIAN

Here, we specify the Hamiltonian used in the lithium hydride ground-state study of Section 4.2. We construct the Hamiltonian at an internuclear distance of 2.5 Å using PySCF [40] and the STO-3G basis set. We consider the LiH molecule oriented with the bond along the x -direction. For this orientation, the unoccupied $2p_y$ and $2p_z$ orbitals are perpendicular to the LiH bond and along with the core orbitals, are assumed frozen; thus, we do not consider the frozen orbitals in the quantum simulations. This Hamiltonian is then transformed into the particle/hole representation [4] and mapped to a qubit Hamiltonian using parity mapping [8, 38], which results in a four-qubit representation after applying qubit tapering to account for symmetries due to spin up and spin down particle number conservation [7]. The full Hamiltonian, written as a linear combination of tensor products of Pauli operators, has 100 terms. See Table 3.

Table 3. LiH Hamiltonian Terms to Six Decimal Places

0.567662 IIII	-0.025425 IZXX	0.013812 ZXXX	-0.011521 XZYY	-0.008083 ZXZX
0.245088 IIZI	0.025425 IZYY	0.013812 IXXX	0.011521 XIYY	-0.008083 IXZX
-0.245088 ZIII	0.025425 XXIZ	-0.013812 ZXYX	0.011521 XZXX	0.008083 ZXIX
-0.190085 IIZZ	-0.025425 YYIZ	-0.013812 IXYX	-0.011521 XIXX	0.008083 IXIX
-0.190085 ZZII	-0.019768 IIXZ	-0.013812 XXZX	0.010474 IIXX	-0.006835 ZXXZ
-0.107219 IZIZ	-0.019768 IIXI	0.013812 YYZX	-0.010474 IYYX	-0.006835 IXXZ
0.101581 IZII	0.019768 XZII	0.013812 XXIX	0.010474 XXII	-0.006835 ZXXI
-0.101581 IIIZ	-0.019768 XIII	-0.013812 YYIX	-0.010474 YYII	-0.006835 IXXI
0.098833 IZZI	-0.018582 XXZI	-0.012909 ZXZI	-0.009880 XZXI	-0.006835 XZZX
0.098833 ZIIZ	0.018582 YYZI	-0.012909 IXZI	0.009880 XIXI	0.006835 XZIX
-0.096556 ZIZI	0.018582 ZIXX	-0.012909 ZIZX	-0.009880 XZXZ	0.006835 XIZX
0.079438 ZZZZ	-0.018582 ZIYY	0.012909 ZIIX	0.009880 XIXZ	-0.006835 XIIX
-0.060240 ZZZI	0.017442 IZZX	-0.011861 XZZI	0.009298 ZZXI	-0.004511 ZXZZ
0.060240 ZIZZ	-0.017442 IZIX	0.011861 XIZI	0.009298 ZZXZ	-0.004511 IXZZ
-0.053253 IZZZ	0.017442 ZXIZ	-0.011861 ZIXZ	-0.009298 XZZZ	0.004511 ZZZX
0.053253 ZZIZ	0.017442 IXIZ	-0.011861 ZIXI	0.009298 XIZZ	-0.004511 ZZIX
0.033028 XXXX	0.016652 IZXZ	-0.011521 XXXZ	-0.009044 IIZX	-0.003631 XXZZ
-0.033028 YYXX	0.016652 IZXI	0.011521 YYXZ	0.009044 IIIX	0.003631 YYZZ
-0.033028 XXYX	0.016652 XZIZ	-0.011521 XXXI	0.009044 ZXII	0.003631 ZZYX
0.033028 YYYY	-0.016652 XIIZ	0.011521 YYXI	0.009044 IXII	-0.003631 ZZXX

ACKNOWLEDGMENTS

The authors thank Bryce Fuller for discussions on the manuscript and Michael Andrews for help on some of the topology concepts.

REFERENCES

- [1] Héctor Abraham et al. 2019. Qiskit: An Open-source Framework for Quantum Computing. Retrieved from <https://doi.org/10.5281/zenodo.2562110>
- [2] Tameem Albash and Daniel A. Lidar. 2018. Adiabatic quantum computation. *Rev. Mod. Phys.* 90, 1 (2018), 015002.
- [3] Eugene L. Allgower and Kurt Georg. 2012. *Numerical Continuation Methods: An Introduction*. Vol. 13. Springer Science & Business Media.
- [4] Panagiotis Kl. Barkoutsos, Jerome F. Gonthier, Igor Sokolov, Nikolaj Moll, Gian Salis, Andreas Fuhrer, Marc Ganzhorn, Daniel J. Egger, Matthias Troyer, Antonio Mezzacapo, Stefan Filipp, and Ivano Tavernelli. 2018. Quantum algorithms for electronic structure calculations: Particle-hole Hamiltonian and optimized wave-function expansions. *Phys. Rev. A* 98, 2 (2018), 022322.
- [5] Dimitri P. Bertsekas. 1999. *Nonlinear Programming* (2nd ed.). Athena Scientific, Belmont, MA.
- [6] Dimitri P. Bertsekas and John N. Tsitsiklis. 2000. Gradient convergence in gradient methods with errors. *SIAM J. Optimiz.* 10, 3 (2000), 627–642.
- [7] Sergey Bravyi, Jay M. Gambetta, Antonio Mezzacapo, and Kristan Temme. 2017. Tapering off qubits to simulate fermionic Hamiltonians. arXiv:1701.08213. Retrieved from <https://arxiv.org/abs/1701.08213>.
- [8] S. Bravyi and A. Kitaev. 2002. Fermionic quantum computation. *Ann. Phys.* 298, 1 (2002), 210–226.
- [9] Marco Cerezo, Akira Sone, Tyler Volkoff, Lukasz Cincio, and Patrick J. Coles. 2021. Cost function dependent barren plateaus in shallow parametrized quantum circuits. *Nat. Commun.* 12, 1 (2021), 1–12.
- [10] Alba Cervera-Lierta, Jakob S. Kottmann, and Alán Aspuru-Guzik. 2021. Meta-variational quantum eigensolver: Learning energy profiles of parameterized hamiltonians for quantum simulation. *P R X Quant.* 2 (2021), 020329.
- [11] Ming-Cheng Chen, Ming Gong, Xiaosi Xu, Xiao Yuan, Jian-Wen Wang, Can Wang, Chong Ying, Jin Lin, Yu Xu, Yulin Wu, et al. 2020. Demonstration of adiabatic variational quantum computing with a superconducting quantum coprocessor. *Phys. Rev. Lett.* 125, 18 (2020), 180501.

- [12] John B. Collins, Paul von R. Schleyer, J. Stephen Binkley, and John A. Pople. 1976. Self-consistent molecular orbital methods. XVII. Geometries and binding energies of second-row molecules. A comparison of three basis sets. *J. Chem. Phys.* 64, 12 (1976), 5142–5151.
- [13] Gavin E. Crooks. 2019. Gradients of parameterized quantum gates using the parameter-shift rule and gate decomposition. arXiv:1905.13311. Retrieved from <https://arxiv.org/abs/1905.13311>.
- [14] Jelena Diakonikolas and Michael I. Jordan. 2019. Generalized momentum-based methods: A Hamiltonian perspective. *SIAM Journal on Optimization* 31, 1 (2021), 915–944.
- [15] Daniel J. Egger, Jakub Marecek, and Stefan Woerner. 2021. Warm-starting quantum optimization. *Quantum* 5 (2021), 479.
- [16] Edward Farhi, Jeffrey Goldstone, and Sam Gutmann. 2014. A quantum approximate optimization algorithm. arXiv:1411.4028. Retrieved from <https://arxiv.org/abs/1411.4028>.
- [17] A. Garcia-Saez and J. I. Latorre. 2018. Addressing hard classical problems with adiabatically assisted variational quantum eigensolvers. arXiv:1806.02287. Retrieved from <https://arxiv.org/abs/1806.02287>.
- [18] Warren J. Hehre, Robert F. Stewart, and John A. Pople. 1969. Self-consistent molecular orbital methods. I. Use of Gaussian expansions of Slater-Type atomic orbitals. *J. Chem. Phys.* 51, 6 (1969), 2657–2664.
- [19] Sabine Jansen, Mary-Beth Ruskai, and Ruedi Seiler. 2007. Bounds for the adiabatic approximation with applications to quantum computation. *J. Math. Phys.* 48, 10 (2007), 102111.
- [20] 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.
- [21] Tosio Kato. 1980. *Perturbation Theory for Linear Operators* (corrected printing of the 2nd ed.). Springer-Verlag, New York.
- [22] Wim Lavrijsen, Ana Tudor, Juliane Müller, Costin Iancu, and Wibe de Jong. 2020. Classical optimizers for noisy intermediate-scale quantum devices. In *Proceedings of the IEEE International Conference on Quantum Computing and Engineering (QCE'20)*. 267–277. <https://doi.org/10.1109/QCE49297.2020.00041>
- [23] Nicolas Loizou and Peter Richtárik. 2020. Momentum and stochastic momentum for stochastic gradient, Newton, proximal point and subspace descent methods. *Comput. Optimiz. Appl.* 77, 3 (2020), 653–710.
- [24] Shunji Matsuura, Samantha Buck, Valentin Senicourt, and Arman Zaribafyan. 2021. Variationally scheduled quantum simulation. *Phys. Rev. A* 103, 5 (May 2021), 052435. <https://doi.org/10.1103/PhysRevA.103.052435>
- [25] Sam McArdle, Tyson Jones, Suguru Endo, Ying Li, Simon C. Benjamin, and Xiao Yuan. 2019. Variational ansatz-based quantum simulation of imaginary time evolution. *NPJ Quant. Inf.* 5, 1 (2019), 1–6.
- [26] Jarrod R. McClean, Sergio Boixo, Vadim N. Smelyanskiy, Ryan Babbush, and Hartmut Neven. 2018. Barren plateaus in quantum neural network training landscapes. *Nat. Commun.* 9, 1 (2018), 1–6.
- [27] Jarrod R. McClean, Jonathan Romero, Ryan Babbush, and Alán Aspuru-Guzik. 2016. The theory of variational hybrid quantum-classical algorithms. *New J. Phys.* 18, 2 (2016), 023023.
- [28] Kosuke Mitarai, Yuya O. Nakagawa, and Wataru Mizukami. 2020. Theory of analytical energy derivatives for the variational quantum eigensolver. *Phys. Rev. Res.* 2, 1 (2020), 013129.
- [29] Kosuke Mitarai, Tennin Yan, and Keisuke Fujii. 2019. Generalization of the output of a variational quantum eigensolver by parameter interpolation with a low-depth ansatz. *Phys. Rev. Appl.* 11, 4 (2019), 044087.
- [30] 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. *Quant. Sci. Technol.* 3, 3 (2018), 030503.
- [31] James R. Munkres. 2000. *Topology* (2nd ed.). Prentice Hall, Upper Saddle River, NJ.
- [32] Ken M. Nakanishi, Keisuke Fujii, and Syngde Todo. 2020. Sequential minimal optimization for quantum-classical hybrid algorithms. *Phys. Rev. Res.* 2, 4 (2020), 043158.
- [33] Giacomo Nannicini. 2019. Performance of hybrid quantum-classical variational heuristics for combinatorial optimization. *Phys. Rev. E* 99, 1 (2019), 013304.
- [34] 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. *Nat. Commun.* 5, 1 (2014), 1–7.
- [35] Boris T. Polyak and Anatoli B. Juditsky. 1992. Acceleration of stochastic approximation by averaging. *SIAM J. Contr. Optimiz.* 30, 4 (1992), 838–855.
- [36] Walter Rudin. 1976. *Principles of Mathematical Analysis* (3rd ed.). McGraw-Hill, New York, NY.
- [37] Maria Schuld, Ville Bergholm, Christian Gogolin, Josh Izaac, and Nathan Killoran. 2019. Evaluating analytic gradients on quantum hardware. *Phys. Rev. A* 99, 3 (2019), 032331.
- [38] Jacob T. Seeley, Martin J. Richard, and Peter J. Love. 2012. The Bravyi-Kitaev transformation for quantum computation of electronic structure. *J. Chem. Phys.* 137, 22 (2012), 224109.

- [39] James C. Spall. 1992. Multivariate stochastic approximation using a simultaneous perturbation gradient approximation. *IEEE Trans. Automat. Contr.* 37, 3 (1992), 332–341.
- [40] Q. M. Sun, T. C. Berkelbach, N. S. Blunt, G. H. Booth, S. Guo, Z. D. Li, J. Z. Liu, J. D. McClain, E. R. Sayfutyarova, S. Sharma, S. Wouters, and G. K. L. Chan. 2018. PySCF: The Python-based simulations of chemistry framework. *Wiley Interdiscipl. Rev. Comput. Molec. Sci.* 8, 1 (2018), e1340.
- [41] Kevin J. Sung, Matthew P. Harrigan, Nicholas C. Rubin, Zhang Jiang, Ryan Babbush, and Jarrod R. McClean. 2020. An exploration of practical optimizers for variational quantum algorithms on superconducting qubit processors. arXiv:2005.11011v1. Retrieved from <https://arxiv.org/abs/2005.11011v1>.
- [42] Stephen J. Wright. 2015. Coordinate descent algorithms. *Math. Program.* 151, 1 (2015), 3–34.
- [43] Xiao Yuan, Suguru Endo, Qi Zhao, Ying Li, and Simon C. Benjamin. 2019. Theory of variational quantum simulation. *Quantum* 3 (2019), 191.

Received February 2021; revised June 2021; accepted August 2021