

Double-bracket quantum algorithms for high-fidelity ground-state preparation

Matteo Robbiati ^{1,2,*}, Edoardo Pedicillo,^{2,3,*} Andrea Pasquale,^{2,3,4,*} Xiaoyue Li ^{5,*}, Oriel Kiss,^{1,6,*} Andrew Wright ⁷, Renato M. S. Farias ^{3,8}, Khanh Uyen Giang ⁵, Jeongrak Son ⁵, Johannes Knörzer,⁹ Siong Thye Goh ¹⁰, Jun Yong Khoo ¹⁰, Nelly H. Y. Ng,⁵ Zoë Holmes,⁷ Stefano Carrazza ^{1,2,3,4} and Marek Gluza ^{5,†}

¹European Organization for Nuclear Research (CERN), Geneva 1211, Switzerland

²TIF Lab, Dipartimento di Fisica, Università degli Studi di Milano, Milan, Italy

³Quantum Research Center, Technology Innovation Institute, Abu Dhabi, United Arab Emirates

⁴INFN Sezione di Milano, Milan, Italy

⁵Nanyang Quantum Hub, School of Physical and Mathematical Sciences, Nanyang Technological University, 21 Nanyang Link, Singapore 637371, Singapore

⁶Department of Nuclear and Particle Physics, University of Geneva, Geneva 1211, Switzerland

⁷Institute of Physics, Ecole Polytechnique Fédérale de Lausanne (EPFL), Lausanne, Switzerland

⁸Instituto de Física, Universidade Federal do Rio de Janeiro, P.O. Box 68528, Rio de Janeiro, Rio de Janeiro 21941-972, Brazil

⁹Institute for Theoretical Studies, ETH Zurich, Zurich 8092, Switzerland

¹⁰Institute of High Performance Computing (IHPC), Agency for Science, Technology and Research (A*STAR), 1 Fusionopolis Way, #16-16 Connexis, Singapore 138632, Singapore



(Received 12 December 2024; accepted 14 May 2026; published 17 July 2026)

Ground-state preparation is a central application for quantum computers but remains challenging in practice. In this work, we quantitatively investigate the performance and gate counts of double-bracket quantum algorithms (DBQAs) for ground-state preparation. We propose a practical strategy in which DBQAs refine initial state preparation circuits, and we compile them for Heisenberg chains using controlled-Z and single-qubit gates. Warm-started DBQAs consistently improve both the energy and ground-state fidelity relative to the initial states provided by variational ansätze, indicating that DBQAs offer an effective unitary synthesis method. To demonstrate compatibility with near-term hardware, we executed a proof-of-concept example on IBM devices. With error mitigation, we observed a statistically significant improvement over the corresponding warm-start circuit. Furthermore, numerical emulations for the same system size indicate that executing DBQAs on Quantinuum's hardware could achieve similar cost-function gains without requiring error mitigation. These findings suggest that DBQAs are a promising approach for enhancing ground-state approximations on near-term quantum devices.

DOI: [10.1103/jz88-32rb](https://doi.org/10.1103/jz88-32rb)

I. INTRODUCTION

Approximating the ground state of a target Hamiltonian is a challenging problem that has been widely investigated in quantum and classical computing. Quantum phase estimation (QPE) [1–5] and quantum signal processing [6,7] are expected to be the go-to methods for computing ground-state energies in the fault-tolerant era, but the circuit depths they require are prohibitive on near-term devices. Instead, in the near term, parametrized circuits are widely used for cost-function optimization [8–10], but lack convergence guarantees and often encounter numerous optimization barriers [11–16].

Recently, double-bracket quantum algorithms (DBQAs) have been introduced for Hamiltonian diagonalization [17],

imaginary-time evolution (ITE) [18,19], and quantum signal processing [20]. These methods are grounded in Riemannian gradient descent [21], which provides a rigorous basis for their performance and convergence guarantees [18]. The DBQA framework has also enabled a derivation of Grover's algorithm directly from optimization principles on Riemannian manifolds [22]. Taken together, these results indicate that DBQAs provide a unifying framework for ground-state preparation. However, to be practically viable, the gate cost must be small. In this article, we investigate the potential of DBQAs for preparing ground states on early fault-tolerant quantum devices [23]. We focus on devices without error correction [24] and therefore quantify the cost in terms of the number of two-qubit gates, which are more difficult to implement than single-qubit gates [25]. However, DBQAs can also be advantageous on fault-tolerant devices, where the cost is typically measured by the number of T gates. Unlike quantum phase estimation [3] or block-encoding formulations of quantum signal processing [7], DBQAs can be implemented as near-term quantum circuits without the need for auxiliary qubits. Crucially, unlike parametrized circuits, DBQAs involve classical optimization only for performance improvement and can be analytically guaranteed to converge [18,21,26,27].

*These authors contributed equally to this work.

†Contact author: marekludwik.gluza@ntu.edu.sg




FIG. 1. We propose a two-stage ground-state preparation protocol: First, apply a relatively short-depth warm-start circuit; second, apply a DBQA circuit to further the ground-state preparation fidelity.

In this work, as schematically shown in Fig. 1, we propose a two-stage ground-state preparation protocol: First, apply an existing state preparation method that uses relatively short-depth circuits; second, apply a DBQA to further improve the state preparation. We selected the variational quantum eigensolver (VQE) as our initialization method not only for its low overhead and ease of unitary retrieval, but more importantly to emphasize that DBQA can refine even limited approximations, such as those achievable with short-depth circuits on near-term hardware. However, we stress that our method is largely agnostic to the choice of initialization method and there exist many potential candidates depending on the problem, including parametrized circuits [28], the Hamiltonian variational ansatz (HVA) [29–31], Hartree-Fock circuits [32], tensor networks [33,34], quantum imaginary-time evolution (QITE) [35–38], coupled cluster [39], dissipative [40], wall-Chebyshev [41], and qubitization [1–3] approaches. The key distinction between double-bracket diagonal-operator iteration (DB-DOI) and the aforementioned early fault-tolerant methods is that DB-DOI provides formal guarantees while relying solely on unitary operations.

To demonstrate the effectiveness of our method, we perform numerical simulations of DBQA for Heisenberg chains. We investigate the dependence of DBQA circuit depths on the quality of the initial approximation of the ground state. As a specific example, for $L = 10$ qubits, a single DBQA step (with circuit depth of 50–100 CZ gates per qubit) can improve the energy estimation by an order of magnitude when initialized through a parametrized circuit with energy closer to the ground state than the first excited state. Finally, we validate our approach with a ten-qubit experiment on IBM’s quantum hardware and perform numerical emulation of Quantinuum’s hardware.

II. DOUBLE-BRACKET DIAGONAL-OPERATOR ITERATIONS

DBQAs are based on differential equations called double-bracket flows (DBFs), which have applications in quantum many-body systems [42–45] and control theory [21]. We discuss two DBFs that play a role in quantum computing. We will summarize the DBQA framework as the conjunction of

two approximations that produce quantum circuits that approximate the ground state of an input Hamiltonian $\hat{H}_0$. The Brockett’s DBF [46] is given by

$$\partial_t \hat{H}_0(\ell) = [[\hat{D}, \hat{H}_0(\ell)], \hat{H}_0(\ell)] \quad (1)$$

and we set $\hat{H}_0(0) = \hat{H}_0$. This equation preserves the spectrum of $\hat{H}_0$ if the operator $\hat{D}$ is Hermitian. If $\hat{D}$ is diagonal and non-degenerate, then $\hat{H}_0(\infty)$ will be diagonal too. Moreover, the spectrum of $\hat{H}_0(\infty)$ will be sorted according to the ordering of the eigenvalues of $\hat{D}$, which is an important feature when using Brockett’s DBF for ground-state preparation [21].

The first approximation to the DBF made in the DBQA framework is a first-order discretization of the differential equation [Eq. (1)]. We set s to be a small Euler-step duration and arrive at a discrete sequence of Hamiltonians defined recursively by

$$\hat{H}_{k+1} = e^{s[\hat{D}, \hat{H}_k]} \hat{H}_k e^{-s[\hat{D}, \hat{H}_k]}. \quad (2)$$

We notice that

$$\hat{H}_{k+1} = \hat{H}_k + s[[\hat{D}, \hat{H}_k], \hat{H}_k] + O(s^2), \quad (3)$$

which links Eqs. (1) and (2). Reference [17] explored the convergence of this sequence to the DBF in the limit $s \rightarrow 0$ and proposed to additionally vary the operator $\hat{D}$ in every step k to reduce the number of steps of the algorithm. Reference [21] details that the sequence $\hat{H}_k$ converges to a diagonal fixed point, similarly as the DBF. This means that

$$\hat{R}_k = e^{-s[\hat{D}, \hat{H}_0]} e^{-s[\hat{D}, \hat{H}_1]} \dots e^{-s[\hat{D}, \hat{H}_k]} \quad (4)$$

converges to a diagonalization unitary. This realizes the first benefit of the DBQA framework: We obtain an analytical pathway to perform unitary synthesis of circuits approximating eigenstates. Let us remark that Brockett’s DBF is a Riemannian gradient flow, which indicates its local optimality and predicts an exponentially fast convergence when initialized in the basin of attraction of the fixed point [21,26,27].

The second approximation step in the DBQA framework is to approximate exponentials of commutators by a product of exponentials. Specifically, we can use that [17]

$$e^{-s[\hat{D}, \hat{H}_k]} = e^{-i\sqrt{s}\hat{H}_k} e^{i\sqrt{s}\hat{D}} e^{i\sqrt{s}\hat{H}_k} e^{-i\sqrt{s}\hat{D}} + O(s^{3/2}). \quad (5)$$

The benefit of this approximation is that each of the four exponentials in the product can be performed on the quantum computer. For example, $\hat{D}$ could be a magnetic field and we assume that $\hat{H}_0$ is a local Hamiltonian. We complete the use of this approximation by redefining the iteration of Hamiltonians $\hat{H}_k$ to be defined by

$$\hat{H}_{k+1} = e^{i\sqrt{s}\hat{D}} e^{-i\sqrt{s}\hat{H}_k} e^{-i\sqrt{s}\hat{D}} \hat{H}_k e^{i\sqrt{s}\hat{D}} e^{i\sqrt{s}\hat{H}_k} e^{-i\sqrt{s}\hat{D}} \quad (6)$$

and note that we have used $e^{i\sqrt{s}\hat{H}_k} \hat{H}_k e^{-i\sqrt{s}\hat{H}_k} = \hat{H}_k$.

To summarize the above, we track the unitary operations that will be applied to the quantum state to obtain an eigenstate approximation. For this, let us denote by $\hat{U}_k$ the unitary that rotates the initial Hamiltonian to the Hamiltonian after the k th iteration step by

$$\hat{H}_k = \hat{U}_k^\dagger \hat{H}_0 \hat{U}_k. \quad (7)$$

We then can use that $e^{it\hat{H}_k} = U_k^\dagger e^{it\hat{H}_0} U_k$. With this, we arrive at the following recursive sequence of unitaries $\hat{U}_k$ that starts with $\hat{U}_0 = \mathbb{1}$ and proceeds by

$$\hat{U}_{k+1} = \hat{U}_k e^{i\sqrt{s}\hat{D}} \hat{U}_k^\dagger e^{i\sqrt{s}\hat{H}_0} \hat{U}_k e^{-i\sqrt{s}\hat{D}}. \quad (8)$$

In the limit $s \rightarrow 0$, we will obtain convergence to Brockett's DBF [17], with the difference to Eq. (1) that the recursive definition involves only evolutions under $\hat{D}$ and $\hat{H}_0$, and we obtain an actionable quantum algorithm. We will call the unitaries in Eq. (8) *double-bracket diagonal-operator iterations*.

Equation (1) pertains to the Hamiltonian and, based on rotating it, it arrives at a unitary that can be applied to computational basis states to obtain approximations of eigenstates. We next consider a related DBF procedure that is instead focused on the input state rather than the Hamiltonian. Given an initial state $|\psi_0\rangle$, we can consider the following DBF:

$$\partial_\tau \Psi(\tau) = [[\Psi(\tau), \hat{H}_0], \Psi(\tau)], \quad (9)$$

where $\Psi(0) = |\psi_0\rangle\langle\psi_0|$. This is again a Brockett's DBF as in Eq. (1), but in the Schrödinger picture because the density matrix of the input state is being rotated and the Hamiltonian plays the role of the constant target matrix $\hat{D}$. If we define imaginary-time evolution by

$$|\psi(\tau)\rangle = \frac{e^{-\tau\hat{H}_0}|\psi_0\rangle}{\|e^{-\tau\hat{H}_0}|\psi_0\rangle\|}, \quad (10)$$

then we have precisely $\Psi(\tau) = |\psi(\tau)\rangle\langle\psi(\tau)|$ [18]. This observation has led to the proposal of another DBQA called double-bracket quantum imaginary-time evolution (DB-QITE) [18], which we will next summarize in comparison to the above DB-DOI discussion.

The first approximation step of the DBQA framework leads us to

$$|\psi(\tau)\rangle = e^{\tau[|\psi_0\rangle\langle\psi_0|, \hat{H}_0]}|\psi_0\rangle + O(\tau^2). \quad (11)$$

This has the important benefit of providing an analytical ansatz for unitary synthesis that implements ITE, which *a priori* involves a nonunitary propagation. Next, after performing the second approximation step of the DBQA framework, i.e., using a product formula approximation to the exponential of a commutator, we arrive at the recursive relation

$$\hat{U}_{k+1} = e^{i\sqrt{s}\hat{H}_0} \hat{U}_k^\dagger e^{i\sqrt{s}|\psi_0\rangle\langle\psi_0|} \hat{U}_k e^{-i\sqrt{s}\hat{H}_0} \hat{U}_k, \quad (12)$$

which is rooted in $\hat{U}_0 = \mathbb{1}$. Reference [18] shows that Eq. (12) allows to reduce the energy achieved by any circuit $\hat{U}_0 \neq \mathbb{1}$ by repeating it three times, which without the DBQA framework appears to be difficult to derive [47].

A. Contrasting DB-DOI with DB-QITE

Both DB-DOI and DB-QITE qualify as double-bracket quantum algorithms, while their distinction lies in formulation: DB-DOI is naturally expressed in the Heisenberg picture, whereas DB-QITE is defined in the Schrödinger picture. The two can be related by observing that $\hat{D}$ and $\hat{H}_0$ in DB-DOI play a role analogous to $\hat{H}_0$ and $|0\rangle\langle 0|$ in DB-QITE. In particular, $\hat{D}$ in DB-DOI and $\hat{H}_0$ in DB-QITE serve as Brockett's target matrices, which define the linear cost function in the Riemannian gradient descent formulation of the double-bracket flow.

In contrast, $\hat{H}_0$ in DB-DOI and $|0\rangle\langle 0|$ in DB-QITE are the dynamical variables that are iteratively transformed toward the desired fixed point.

A priori, the fixed point of DB-DOI does not directly address ground-state preparation. However, Brockett's DBFs exhibit a sorting property: Generically, the spectrum of the limiting operator H_∞ is ordered according to the eigenvalue ordering encoded in $\hat{D}$. Consequently, the ground state of $\hat{D}$ corresponds to the ground state of $\hat{H}_0$ upon convergence of DB-DOI. In contrast, the fixed point of DB-QITE is explicitly the ground state itself. Both algorithms originate from Brockett's DBF, a Riemannian steepest-descent flow [48]. For discussions in the context of quantum computing, see Refs. [18–20,22], and for a general treatment of DBFs, see the monograph [21]. The quantum algorithmic prescription for both DB-DOI [Eq. (8)] and DB-QITE [Eq. (12)] is a recursive unitary synthesis.

The input Hamiltonian $\hat{H}_0$ is the variable in DB-DOI and is typically full rank. In contrast, for DB-QITE the variable $|0\rangle\langle 0|$ is rank 1. Thus, this makes Brockett's DBF sorting property far more striking: instead of sorting the full spectrum, we only shift the energy of the pure state to match the ground state. On the other hand, it can be beneficial in other cases that DB-DOI can also converge to excited eigenstates by initializing in an excited state of $\hat{D}$. To achieve that, DB-QITE would need to perform typical scanning of the folding spectrum, i.e., consider $\hat{H}'_0(\lambda) = (\hat{H}_0 - \lambda)^2$ with various λ because then the ground state of $\hat{H}'_0(\lambda)$ will be the eigenstate of $\hat{H}_0$ with eigenvalue closest to λ .

It is no coincidence that DB-QITE involves $e^{-i\sqrt{s}|0\rangle\langle 0|}$, where a similar term appears in Grover's algorithm: The latter can be derived by the two steps of the DBQA framework [22]. This type of operation has recently led to a unitary synthesis for quantum signal processing that does not involve postselection. For both these insights, the rank-1 property of the variable has been key. It might be that when the DBQA involves $\hat{D}$, which is full rank, it is harder to uncover new properties.

It is crucial to note that it is possible to express $e^{-i\sqrt{s}|0\rangle\langle 0|}$ in terms of CZ and single-qubit gates with a linear scaling in the number of qubits. This holds even for architectures with qubits placed in a single row [49–51]; see Ref. [52] for a proposal of implementing such operations natively in trapped ions. Nevertheless, for contemporary prototypes of quantum computers, the overhead in DB-QITE can still pose challenges. Instead, we will show that DB-DOI can be effective on noisy hardware by having $\hat{D}$ involve only single-qubit terms. Additionally, as a rule of thumb for noisy quantum hardware, the first step of DB-QITE roughly involves three queries to Hamiltonian simulation, while DB-DOI involves only one query. For this reason, DB-DOI is preferred over DB-QITE when implementing the DBQA approach on quantum hardware.

On the other hand, we will see that a single DB-QITE step tends to reduce the energy by a larger amount than DB-DOI, which instead requires additional iterations to achieve a comparable energy reduction. For both DBQAs, we present numerical simulations for up to $L = 20$ qubits.

Finally, we remark that we can use the quantum dynamic programming [53] approach for executing the recursion (12)

with polynomial circuit depth but with an exponential circuit width as a trade-off. While the same can be done for DB-DOI, it would require coding the diagonal operator $\hat{D}$ in a quantum state.

B. Warm starts for DB-DOI

Reference [17] introduced DB-DOI for eigenstate preparation, building on known convergence results [21,27], but it did not address how to specifically target ground states. In the present work, we demonstrate that warm starts significantly enhance ground-state preparation, particularly when low gate counts are a priority. In the subsequent Ref. [18], DB-QITE was developed and also leverages warm starts. We now discuss how warm starts function in each of these methods, highlighting both the practical benefits observed here and directions for future optimization.

As we explained above, the sorting property of Brockett's DBF implies that as long as the initial state DB-DOI is the ground state of $\hat{D}$, then DB-DOI will rotate it to the ground state of $\hat{H}_0$. To estimate the number of steps required, we qualitatively consider two key factors: the spectral gap of the Hamiltonian and the exponential convergence behavior of DBQAs. Once DB-DOI has advanced sufficiently, the operator $\hat{H}_k$ will be largely diagonal, and consequently the bracket $[\hat{D}, \hat{H}_k]$ becomes small. By "small," we mean that the resulting unitaries $\hat{R}_k$ cannot change the energy by more than the spectral gap in a single step. Consequently, although the algorithm converges at the exponential rate characteristic of approximate gradient descent, the limited energy change per step naturally drives the state toward the closest eigenstate in energy. When DB-DOI is initialized with a warm start that is sufficiently close to the ground state, then the basin of attraction will indeed contain the ground state.

The concept of a warm start for DB-QITE arises naturally. Consider a low-energy state $|\psi_0\rangle = \hat{U}_0|0\rangle$, where $\hat{U}_0$ prepares the initial state from the reference state $|0\rangle$. In this case, one can take $\hat{U}_0 \neq \mathbb{1}$ as the starting point, and the DB-QITE recursion then proceeds as

$$\hat{U}_{k+1} = e^{i\sqrt{s}\hat{H}} \hat{U}_k^\dagger e^{-i\sqrt{s}|0\rangle\langle 0|} \hat{U}_k e^{-i\sqrt{s}\hat{H}} \hat{U}_k. \quad (13)$$

Reference [18] proved that in each step, the energy is improved. This is essentially the same as Eq. (12) but with a different basic reflection operation, and it is as if we performed some $\hat{U}_0$ in place of the first DB-QITE step.

We next discuss warm starting in DB-DOI, where we will work in the Heisenberg picture instead of the Schrodinger picture in the case of DB-QITE. Set

$$\hat{H}_{\text{ws}} = \hat{U}_0^\dagger \hat{H}_0 \hat{U}_0 \quad (14)$$

and for $k > 1$ we consider the rotated Hamiltonians in Eq. (7) where $\hat{U}_k$ is defined by Eq. (8) rooted in the warm start $\hat{U}_0$. As we will see in Sec. II C, one does not need to construct U_0 to have specific diagonalization properties. Instead, for DB-DOI, the warm-starting mechanism can be effective by relying on methods that yield circuits $\hat{U}_0$ for approximating the ground state of $\hat{H}_0$.

We apply the DB-DOI unitary operations to the reference state $|\psi_k\rangle := \hat{U}_k|0\rangle$. Then, we work in the Heisenberg picture

and recover the correct energy by setting

$$\tilde{E}_k := \langle \psi_k | \hat{H}_0 | \psi_k \rangle = \langle 0 | \hat{H}_k | 0 \rangle. \quad (15)$$

The latter value converges to an eigenvalue by virtue of the diagonalization property of Brockett's DBF, and consequently $|\psi_k\rangle$ serves as an approximation to the corresponding eigenstate that can be prepared on a quantum computer. In practice, the initial steps of the optimization are the most significant. We now turn to the explicit form of the DB-DOI unitaries.

C. Explicit unfolding of DB-DOI with warm start

To summarize the above discussions, let us track a few steps of DB-DOI after a warm start. For the first update of the Hamiltonian, we use Eq. (14) so that the energy expectation value agrees with the loss function of the warm-start state $|\psi_0\rangle = U_0|0\rangle$ as follows:

$$\langle 0 | \hat{H}_{\text{ws}} | 0 \rangle = \langle \psi_0 | \hat{H}_0 | \psi_0 \rangle. \quad (16)$$

This exemplifies the difference between the Heisenberg and Schrödinger pictures: The left-hand side of the equality is in the former and the right-hand side is in the latter.

Recall that in the warm-starting convention we set $\hat{U}_0 \neq \mathbb{1}$. In the second step, we will set $r = \sqrt{s}$ and

$$\hat{U}_1 = \hat{U}_0 e^{ir\hat{D}} \hat{U}_0^\dagger e^{ir\hat{H}_0} \hat{U}_0 e^{-ir\hat{D}}. \quad (17)$$

With this, we find $\hat{H}_1 = \hat{U}_1^\dagger \hat{H}_0 \hat{U}_1$. This means that the state will be

$$\langle 0 | \hat{H}_1 | 0 \rangle = \langle \psi_1 | \hat{H}_0 | \psi_1 \rangle, \quad (18)$$

where $|\psi_1\rangle = \hat{U}_1|0\rangle$. Notice again how in DB-DOI we work in the Heisenberg picture, and the reference state does not change, but rather the Hamiltonian is being rotated. We could also rewrite $\hat{U}_1$ in Eq. (17) as

$$\hat{U}_1 = \hat{U}_0 e^{ir\hat{D}} e^{ir\hat{H}_{\text{ws}}} e^{-ir\hat{D}}, \quad (19)$$

but it would be only useful if it were easier to execute evolutions $e^{ir\hat{H}_{\text{ws}}}$ than the composed evolution $\hat{U}_0^\dagger e^{ir\hat{H}_0} \hat{U}_0$.

It is important to notice that $\hat{U}_0$ in Eq. (17) acts after the group commutator approximation; i.e., the DB-DOI step acts before the warm start. This suggests that the DB-DOI step has the role of entangling the computational basis state subtly and preparing it for the action of the warm-start unitary. However, also note that

$$e^{-ir\hat{D}}|0\rangle = |0\rangle \quad (20)$$

whenever $\hat{D}$ is traceless and otherwise we acquire an immaterial global phase. Thus, more explicitly, we have

$$|\psi_1\rangle = \hat{U}_0 e^{ir\hat{D}} \hat{U}_0^\dagger e^{ir\hat{H}_0} \hat{U}_0|0\rangle = \hat{U}_0 e^{ir\hat{D}} \hat{U}_0^\dagger e^{ir\hat{H}_0} |\psi_0\rangle, \quad (21)$$

which means that the Hamiltonian evolution in the first step can indeed be interpreted as acting on the warm-start state $|\psi_0\rangle$. We remark that it is not apparent *a priori* that such sequencing of unitary operations would lower the energy of the state—the DBQA framework provides a sound theoretical grounding for that.

To advance by one more DB-DOI step, we set

$$\hat{U}_2 = \hat{U}_1 e^{ir\hat{D}} \hat{U}_1^\dagger e^{ir\hat{H}_0} \hat{U}_1 e^{-ir\hat{D}} \quad (22)$$

$$= \hat{U}_0 e^{ir\hat{D}} \hat{U}_0^\dagger e^{ir\hat{H}_0} \hat{U}_0 e^{ir\hat{D}} \hat{U}_0^\dagger e^{-ir\hat{H}_0} \hat{U}_0 e^{-ir\hat{D}} \hat{U}_0^\dagger \\ \times e^{ir\hat{H}_0} \hat{U}_0 e^{ir\hat{D}} \hat{U}_0^\dagger e^{ir\hat{H}_0} \hat{U}_0 e^{-i2r\hat{D}}. \quad (23)$$

Here, the point is that there are in total four queries to the evolution governed by the input Hamiltonian, nine queries to the warm start, and five queries to evolutions governed by $\hat{D}$. The latter count is lowered by fusing unitaries of the same type, e.g., obtaining the trailing unitary $e^{-i2r\hat{D}}$. This leads to

$$|\psi_2\rangle = \hat{U}_2|\mathbf{0}\rangle \quad (24)$$

and again by disregarding a global phase, the implementation of the trailing diagonal unitary $e^{-i2r\hat{D}}$ can be omitted in $\hat{U}_2$. We will say that Eq. (23) is the unfolded form of the second DB-DOI step. To execute multiple steps of DB-DOI on quantum hardware, we first repeat such unfolding of the recursion for further steps until all unitaries appearing in the sequence are generated by either $\hat{D}$ or $\hat{H}_0$.

D. Higher-order product formulas in DBQAs

An essential step of the DBQA formalism is to specify an explicit strategy for compiling the unitaries $\hat{R}_k$ in Eq. (4) as quantum circuits. The most basic approach is to use the first-order group commutator formula in Eq. (5). However, we can replace $\hat{R}_k$ by a higher-order approximation. For example, using Eq. (8) of Ref. [54] we can use invariance as was done in Eq. (6) and obtain the following higher-order DB-DOI recursion:

$$\hat{U}_{k+1} = \hat{U}_k e^{ir\phi\hat{D}} e^{ir\hat{H}_k} e^{-ir(\phi+1)\hat{D}} e^{-ir(1-\phi)\hat{H}_k} e^{ir\hat{D}}, \quad (25)$$

with $\phi = \frac{1}{2}(\sqrt{5} - 1)$ and $r = \sqrt{s}$. This gives $\hat{U}_{k+1}^\dagger \hat{H}_0 \hat{U}_{k+1} \approx e^{s_k[\hat{D}, \hat{H}_k]} \hat{U}_k^\dagger \hat{H}_0 \hat{U}_k e^{-s_k[\hat{D}, \hat{H}_k]}$, with an error bounded as $O(s_k^2)$, and this could be beneficial when the exponential of the commutator can give a large energy gain and therefore should be approximated accurately. In that case, other group commutator approximations could be considered [55–59]. Whenever we use a product formula, we will use the identity $e^{ir_k\hat{H}_k} = \hat{U}_k^\dagger e^{ir_k\hat{H}_0} \hat{U}_k$ and call Hamiltonian simulation [60,61] for each appearance of an evolution operator in $\hat{U}_k$.

This recursive unfolding at each step k results in a circuit depth that grows exponentially with k . More precisely, if a product formula involves n evolution operations, then k steps require $(2n+1)^k$ queries to Hamiltonian simulation. For instance, in Eq. (25) we have $n=2$, so DB-DOI entails 5^k queries, whereas the minimal first-order formula in Eq. (8) requires only 3^k queries. Although the latter demands fewer queries, it may yield smaller energy improvements. At present, we lack evidence that one approach is unambiguously superior. On very noisy hardware, the minimal first-order formula in Eq. (8) is preferable, whereas in other settings it may be advantageous to employ the higher-order formula of Eq. (25), trading fewer recursion steps for potentially larger energy gains per step.

E. Variationally optimizing steps of DB-DOI

We next notice that in each step k of DB-DOI, we may select a different step duration s_k . To study this further, we define for each k

$$\hat{H}_k(s) = e^{s[\hat{D}, \hat{H}_k]} \hat{H}_k(0) e^{-s[\hat{D}, \hat{H}_k]} \quad (26)$$

as a function of a variable s and call it a double-bracket rotation since it satisfies a Heisenberg equation involving two, not one, brackets:

$$\partial_s \hat{H}_k(s) = [[\hat{D}, \hat{H}_k(0)], \hat{H}_k(s)]. \quad (27)$$

This allows us to define a greedy optimization scheme of a cost function $f : \mathbb{C}^{2^L \times 2^L} \rightarrow \mathbb{R}_{\geq}$ by considering a global minimum of how much a double-bracket rotation can reduce the cost function:

$$s_k = \operatorname{argmin}_{s \in \mathbb{R}} f(\hat{H}_k(s)). \quad (28)$$

Given this optimizer, we then set $\hat{H}_{k+1} = \hat{H}_k(s_k)$.

Moreover, we have the freedom to optimize in every step the operator $\hat{D}$. For each DB-DOI step, we parametrize $\hat{D}_k$ as a classical Ising model with nearest-neighbor (NN) interactions:

$$\hat{D}_k(B^{(k)}, J^{(k)}) = \sum_{i=1}^L (B_i^{(k)} \hat{Z}_i + J_i^{(k)} \hat{Z}_{i+1} \hat{Z}_i), \quad (29)$$

where L is the number of sites in the chain, the parametrization is implemented through the coefficients $\{B_i\}_{i=1}^L$ and $\{J_i\}_{i=1}^L$, and the superscript k highlights the optimization procedure is repeated for each DBQA step (see Appendix C). We refer the reader to Ref. [62] for a detailed discussion on strategies for optimizing the D_k operators. As we will mention also explicitly below, we can achieve sizable energy gains with DB-DOI involving the constant magnetic field $\hat{D}(1, 0) \equiv \sum_{i=1}^L \hat{Z}_i$. However, for that step duration, the optimization is still essential.

III. NUMERICAL RESULTS FOR XXZ

Consider the one-dimensional XXZ Heisenberg model,

$$\hat{H}_{\text{XXZ}} = \sum_{i=1}^L (\hat{X}_i \hat{X}_{i+1} + \hat{Y}_i \hat{Y}_{i+1} + \Delta \hat{Z}_i \hat{Z}_{i+1}), \quad (30)$$

where the subscript i indicates that the Pauli operators act on the i th qubit, and Δ is the anisotropy parameter that decides whether the system is gapped or gapless in the thermodynamic limit. The XXZ model [Eq. (30)] is well understood by means of the Bethe ansatz [63–65] and tensor networks [66]. In Sec. III A, we study DB-DOI with warm start by VQE [67] and use periodic boundary conditions [note the summation until L in Eq. (30), and we set $X_{L+1} \equiv X_1$, etc.]. On hardware, it is often more natural to consider open-boundary conditions ($X_{L+1} \equiv 0$, etc.), which we shall also use when studying larger system sizes.

Our objective is to achieve high fidelity for the ground-state preparation task, which we define qualitatively in light of the following considerations. For quantum many-body systems, perturbation theory is inherently unstable, as typical vectors in a high-dimensional Hilbert space exhibit exponentially small overlaps; this phenomenon is sometimes referred to as the orthogonality catastrophe [68]. Consequently, fidelity with

the ground state should be regarded as an especially sensitive measure of success. In practice, physics often addresses this sensitivity by shifting from strictly quantitative to more qualitative criteria: Instead of demanding high fidelity, one may seek a state of sufficiently low energy, and instead of requiring exact ground-state preparation, one may settle for a low-temperature state. In this work, we are guided by the notion of ground-state fidelity witnesses [69], and we refer to a state $|\psi\rangle$ as a high-fidelity ground-state preparation whenever its energy expectation value $\langle \hat{H}_0 \rangle_\psi$ lies substantially below the first excited-state energy E_1 . When this is true, then the preparation has energy close to the ground-state energy E_0 , and so the fidelity witness $\mathcal{F}$ defined as

$$\mathcal{F} := 1 - (\langle \hat{H}_0 \rangle_\psi - E_0) / (E_1 - E_0) \quad (31)$$

will have a nontrivial value $\mathcal{F} \geq 0$. Given that $s\mathcal{F}$ is a lower bound to the fidelity of the preparation, we see intuitively that reaching energy falling within the spectral gap of the model $E_0 \leq \langle \hat{H}_0 \rangle_\psi \leq E_1$ implies a nontrivial fidelity to the ground state. We will say that a preparation reached ground state with high fidelity when the energy expectation value is below the middle of the spectral gap:

$$\langle \hat{H}_0 \rangle_\psi \leq (E_0 + E_1)/2. \quad (32)$$

The choice of setting it in the middle is arbitrary, but it is both permissive enough given the orthogonality catastrophe and stringent enough because all such states will have the actual fidelity $F \geq 50\%$. Below, we will see that warm starts by parametrized circuits can bring the energy expectation value to within the spectral gap. On the other hand, reaching a high-fidelity regime in this qualitative sense becomes difficult for about a dozen qubits. Fortunately, DB-DOI can amplify the fidelity of its warm start. Moreover, we will also consider a problem-specific variational ansatz that performs better. With this, we show that DB-DOI can consistently improve it to arrive in the high-fidelity regime as defined by Eq. (32).

A. Exploring numerically simulated variational optimization of parametrized circuits

The variational quantum eigensolver [28] is a variational quantum computing routine that minimizes the same cost function by optimizing a parametric quantum circuit $\hat{U}_\theta$, i.e., varies parameters θ to find an optimal set of parameters

$$\theta^* = \operatorname{argmin}_\theta \{ \langle 0 | \hat{U}_\theta^\dagger \hat{H}_0 \hat{U}_\theta | 0 \rangle \} \quad (33)$$

that minimizes the energy. Although VQE is known for its variational character, trainability and expressibility problems appear [9,11,12,15,16,70], limiting its effectiveness. However, for DBQAs, we only need an approximation of the ground state and VQE might be effective for this despite its limitations. Thus, we will use VQE as a warm start for DBQA and set $\hat{U}_0 = \hat{U}_{\theta^*}$ in Eq. (14).

As a remark, in this work, we do not use the variational aspect as an optimization loop involving a quantum computer. Since we only need a sufficiently good initial state, we will instead numerically simulate the VQE training process. The optimal set of parameters θ^* obtained from simulation can then be used to initialize the parametrized quantum circuit on quantum hardware as a warm start for DB-DOI. We quanti-




FIG. 2. Hamming-weight-preserving architecture for $L = 4$ qubits and $S = 2$. A single circuit layer consists of RBS gates [see Eq. (34)] connecting nearest and next nearest neighbors. We account for periodic boundary conditions by adding RBS gates connecting the last and first qubits. After a chosen number of layers, we add measurements in the computational basis.

tatively assess whether this approach is viable for DB-DOI. As a general remark, even for $L \gg 10$ qubits where VQE training becomes ineffective, DB-DOI could benefit from a VQE warm start by subdividing $L = KL'$ into K blocks and then initializing DB-DOI from a product state of K VQEs computed by a classical variational loop that we will present below for only one block.

We will report below quantitative results for preparing the ground state of the XXZ model, which satisfies a total-spin conservation [71]. For an even number of qubits L —as in the examples examined in this work—and total spin S , the ground state of the XXZ Hamiltonian is constrained to the half-filling subspace, i.e., a superposition of states associated with total spin $S = L/2$. We account for this symmetry by using a *Hamming-weight-preserving* VQE ansatz [15,16,72–82]. This reduces the search space to a Hilbert space of size $\binom{L}{L/2}$, providing a polynomial compression. Even though this compression is not ideal, i.e., exponential, fulfilling the system symmetries allows for faster training due to circuit updates being constrained to a subspace of interest [79,83–87]. One way to construct a Hamming-weight-preserving ansatz consists of creating a network of Reconfigurable Beam-Splitter (RBS) gates:

$$\begin{array}{c} \text{---} \bullet \text{---} \\ | \quad \theta \\ \text{---} \bullet \text{---} \end{array} = \begin{bmatrix} 1 & 0 & 0 & 0 \\ 0 & \cos \theta & \sin \theta & 0 \\ 0 & -\sin \theta & \cos \theta & 0 \\ 0 & 0 & 0 & 1 \end{bmatrix}. \quad (34)$$

Such a gate is a common building block for constructing circuits that preserve the Hamming weight of bitstrings that label computational basis states in a given superposition. In Fig. 2, we describe the circuit architecture implemented in our numerical experiments. For completeness, we also report in Appendix B results obtained with a hardware-efficient ansatz [88], which explores the full Hilbert space of dimension 2^L . The enhanced expressibility increases the difficulty of training, yet combining with DBQA once again proves beneficial. While employing a symmetry-aware ansatz improves convergence, it does not directly influence the performance of DBQA itself when the fidelity of the warm start is held fixed.

1. Results

After training the VQE circuit using gradient descent with the parameter-shift rule [90] and the ADAM optimizer [91]




FIG. 3. Visualization of the impact of DB-DOI on cost function for a single VQE random seed; see Table I for statistical analysis. (left) Training of VQE (blue lines) for three, four, and five layers of a Hamming-weight-preserving ansatz (hues of blue) achieved ground-state energy residue $\Delta E \approx 1\%$ within 500 training epochs. For more epochs, the initially rapid decrease in the cost function saturates and shows marginal improvement afterward. We initialize DBQA with VQE for selected epochs $\in \{100, 200, 500, 1000, 2000\}$ and optimize DBQA parameters with CMA-ES [89]. In the bottom panel, we show the relative difference value between the achieved energy $\tilde{E}_0$ and the true ground-state energy E_0 . (right) Token cost estimates of DB-DOI by counting the total number of CZ gates required for the complete protocol, which includes the training (using the parameter-shift rule) of the VQE until a target epoch and the optimization of DBQA. The depth of the individual circuits is instead reported in the legend.

until a target epoch, we use it as a warm start for DB-DOI. For each DB-DOI step, we parametrize $\hat{D}_k$ as the classical Ising model with NN interactions as in Eq. (29) and optimize the parameters using the CMA genetic algorithm [89], which does not require explicit computation of the gradients. This parametrization allows us to compile $e^{-it\hat{D}_k}$ with at most two layers of CZ gates, and we use two steps of first-order Trotter-Suzuki decomposition [60] for a short-depth compilation of $e^{-it\hat{H}_{\text{XXZ}}}$ (see Appendix D). A detailed description of the procedure by which we obtained the results we present in this section can be found in Appendix C.

To visualize the advantages of DB-DOI with VQE warm starts, in Fig. 3 we show the results for $L = 10$ qubits obtained by executing the procedure once. We find that DBQA halves

the energy residue if applied to the early training epochs, and essentially reaches the ground state when executed later in the process. More specifically, if we use three layers of the ansatz from Fig. 2, so a warm-start circuit with depth of 12 CZ gates per qubit, then following up with one DB-DOI step yields a circuit with depth 75 CZ gates per qubit. We quantify its performance through the energy approximation ratio

$$\Delta E := |(\tilde{E}_0 - E_0)/E_0|, \quad (35)$$

where $\tilde{E}_0$ is the energy of DB-DOI and E_0 is the true ground-state energy. Figure 3 shows an improvement by an order of magnitude from $\Delta E \approx 1\%$ to $\Delta E \approx 0.1\%$. This cost-function gain is statistically significant; see Table I where we consider 50 executions of DB-DOI for each VQE circuit configura-

TABLE I. (above) Relative energy difference for the XXZ model of Eq. (30) with $L = 10$ qubits, and $\Delta = 0.5$. The estimates with their uncertainties were calculated using the median and the median absolute deviation of a sample of results obtained by repeating the execution 50 times with different initial conditions. (below) Circuit depth expressed as the number of CZ gates per qubit, alongside the cumulative number of CZ gates used to reach $\tilde{E}_0$ (see Appendix C). Warm-started VQE approximations (500 epochs of training) are presented alongside DB-DOI results, executed considering both compiled (DB-DOI) and theoretical double-bracket iterations (DBI) approaches defined in Ref. [17]. For commutator exponentials of DBI, we compute $\hat{R}_k$ defined in Eq. (4) through dense matrix representation, so there is no gate count. Longer VQE training (2000 epochs) is reported in the last column of the table.

Layers	Warm start $1 - \tilde{E}_0/E_0$		1 DB-DOI step $1 - \tilde{E}_0/E_0$		2 DB-DOI steps $1 - \tilde{E}_0/E_0$		1 DBI step $1 - \tilde{E}_0/E_0$	Long VQE training $1 - \tilde{E}_0/E_0$	
3	0.012 ± 0.004		0.0011 ± 0.0007		0.0005 ± 0.0004		0.0009 ± 0.0006	0.010 ± 0.004	
4	0.008 ± 0.004		0.0006 ± 0.0005		0.0002 ± 0.0002		0.0005 ± 0.0004	0.004 ± 0.002	
5	0.005 ± 0.003		0.0003 ± 0.0002		0.0001 ± 0.0001		0.0002 ± 0.0002	0.003 ± 0.002	
	Depth	Cumulative cost	Depth	Cumulative cost	Depth	Cumulative cost	—	Depth	Cumulative cost
3	12	1.44×10^7	75	1.49×10^7	390	1.69×10^7	—	12	5.76×10^7
4	16	2.56×10^7	95	2.62×10^7	490	2.88×10^7	—	16	10.24×10^7
5	20	4.0×10^7	115	4.07×10^7	590	4.38×10^7	—	20	16.0×10^7

TABLE II. Fidelity lower bounds [69] (see Appendix C) extending results presented in Table I.

Layers	Warm start	$k = 1$	$k = 2$	$k = 3$
3	0.83 ± 0.06	0.95 ± 0.01	0.993 ± 0.006	0.997 ± 0.003
4	0.89 ± 0.05	0.992 ± 0.007	0.997 ± 0.003	0.998 ± 0.001
5	0.93 ± 0.04	0.996 ± 0.003	0.998 ± 0.002	0.9992 ± 0.0008

tion. Consistently, DBQA steps improve energy estimation for all depths analyzed. Energy measurements can be translated to fidelity lower bounds [69] as described in Appendix C. Table II shows that one step of DB-DOI allowed us to get $\mathcal{F} \geq 99.6\% \pm 0.3\%$ fidelity.

By increasing the circuit depth of the chosen ansatz of Fig. 2 to 20 CZ gates per qubit, it can be trained to reach $\Delta E \approx 0.4\%$. In general, whenever an initialization method is improved to match the previously top-performing DB-DOI circuit, DBQA can be interfaced with that enhanced initialization. Indeed, when initialized with the 20 CZ gates per qubit circuit depth with $\Delta E \approx 0.4\%$, just one DBQA step again gives an order of magnitude gain $\Delta E \approx 0.03\%$. We were not able to reach $\Delta E \approx 0.03\%$ by training deeper VQE circuits.

The DB-DOI circuits executed to produce Fig. 3 involve a VQE warm start with three layers of the ansatz from Fig. 2, which translates to circuit depth of 12 CZ gates per qubit, and then one DB-DOI step yields depth 75 while two steps have depth 390. The right plot in Fig. 3 shows the training cost of the algorithm quantified in terms of the cumulative number of CZ gates during optimization, which captures the run-time of the classical optimization. More details about the cumulative cost metric can be found in Appendix C. We find that the cost of training is dominated by VQE when running a few DBQA steps. For tasks requiring high ground-state preparation fidelity, VQE and DBQA should be used in sequence, as individually each would necessitate unnecessarily large token expenditures, i.e., cost of execution on commercial hardware. This training cost reduction is summarized in Table I, where we report the gate count in the case where DB-DOI has been compiled into circuits via the DB-DOI formalism (see Appendix D).

Alongside compiled DB-DOI circuits, in the next-to-last column of Table I, we show DB-DOI results using DBI unitaries $e^{s_k[\hat{D}_k, \hat{H}_k]}$ that do not involve the product formula approximations and find that they perform similarly to DB-DOI, indirectly highlighting the usefulness of the higher-order approximation in Eq. (25).

Further investigations are presented in Appendices A, B and E, where we consider (1) one-dimensional XXZ targets of sizes $L \in \{4, 6, 8, 10, 12\}$, (2) a VQE architecture that does not respect XXZ symmetries and so requires longer training, but yields better warm-start energy in shorter depth, and (3) next-nearest-neighbor interactions in the target Hamiltonian. In all these cases studied, we verify that DB-DOI remains advantageous.

B. Exploring numerically simulated Hamiltonian variational ansatz

The Hamiltonian variational ansatz is well suited for the XXZ model. In particular, we will consider up to $L = 20$

qubits with open-boundary conditions and $\Delta = 1$, which leads to the gapless model XXX in the thermodynamic limit. To begin, we use the simple observation that the XXX model is a sum of singlet projections and thus for XXX on a line, the tensor product of singlet states

$$|\text{Singlet}\rangle = 2^{-L/4}(|10\rangle - |01\rangle)^{\otimes L/2} \quad (36)$$

is in the ground state of half of the terms in the Hamiltonian. We can further improve on it using the problem-specific Hamiltonian Variational Ansatz [29–31]

$$|\text{HVA}_p(\mathbf{t})\rangle = \prod_{j=1}^p (e^{-it_{j,e}\hat{H}_e} e^{-it_{j,o}\hat{H}_o}) |\text{Singlet}\rangle, \quad (37)$$

where $\hat{H}_e, \hat{H}_o$ are obtained from $\hat{H}_0$ by restricting to summation over odd (even) indices i such that all terms in $\hat{H}_e, \hat{H}_o$ commute. To understand this ansatz, notice that for $p = 1$,

$$|\text{HVA}\rangle = e^{-it_0\hat{H}_e} e^{-it_1\hat{H}_o} |\text{Singlet}\rangle = U_{\text{HVA}} |\text{Singlet}\rangle. \quad (38)$$

The first unitary acting on the singlets on the even sublattice has the action of creating singlets on the odd sublattice. Thus, the layers can be interpreted as creating a balance between the location of the singlets and putting the singlet sublattices into a superposition. Reference [18] demonstrated that HVA works well as a warm start for DB-QITE, and below we will assess its use for warm-starting DB-DOI.

1. Results

We observe that the HVA with $p = 1, 2, \dots, 6$ can yield excellent approximations to the XXX ground state. When it performs well, it has the advantage of a relatively low gate count. However, it might be difficult to improve the quality of the warm start, as illustrated in the left inset of Fig. 4: When training variational circuits with the COBYLA optimizer—even with randomized basin hopping—the objective function does not always decrease monotonically with increasing number of layers. This “hit-and-miss” behavior may reflect the difficulty of HVA training, similarly to the difficulties due to barren plateaus in VQEs [92–94]. Furthermore, while the HVA provides an excellent approximation for the XXX ground state, its performance is not guaranteed for other systems. In such cases, or when further circuit improvement is needed, DBQAs offer a viable alternative. We find that for all values of p considered, DB-DOI consistently improves upon the warm-start HVA energy; see main panel of Fig. 4.

Finally, the right inset in Fig. 4 shows that DB-DOI achieves performance comparable to DB-QITE. DB-DOI has the advantage of a very inexpensive initial step relative to DB-QITE, though each step reduces the energy by a smaller amount. This limitation can be mitigated: Starting from the same HVA warm start with $p = 1$, similar energies can be




FIG. 4. Performance of DB-DOI applied to the XXZ model with $\Delta = 1$ for $L = 20$ qubits. The main panel shows the energy as a function of the number of CZ gates per circuit per qubit, with points corresponding to $k = 0, 1, 2, 3$ DB-DOI steps (different colors indicate different HVA ansatz depths p for $k = 0$). DB-DOI is warm started from HVA states of varying depth (left inset), where the colored circles mark the DB-DOI initial energies in the main panel. In the left inset, we show the relative energy to the ground-state energy λ_0 to visualize the system size scaling of HVA for $L = 12, 16, 20$ qubits. To explore HVA for varying system sizes $L = 12, 16, 20$, we offset the changing target energy by plotting the HVA energy relative to the exact ground state ($\Delta E_{HVA} = E_{HVA} - E_0$) as a function layer number p that suggests that for $p = 1$ HVA is limited by expressivity while for $p \geq 3$ it is possible to find improved ground-state approximations but training becomes a limitation. The right inset compares DB-DOI and DB-QITE for the same initialization, showing that DB-DOI achieves similar energy reductions as DB-QITE with a few hundred CZ gates per qubit. Note that the number of CZ gates plotted refers to the circuit depth and does not include the training's cost.

obtained by either DB-DOI with $k = 3$ steps or DB-QITE with $k = 2$ steps, both having comparable CZ gate counts. It is important to note that the DB-QITE gate counts assume all-to-all connectivity; on fixed architectures without adaptive use of auxiliary qubits and measurements, the actual gate count would be higher than reported.

IV. EXECUTION OF DB-DOI ON IBM QUANTUM HARDWARE

Using remote access, we performed a proof-of-concept experiment on an IBM superconducting quantum chip for the $L = 10$ qubit XXZ model with open-boundary conditions. The initial state is obtained by training two layers of nearest-neighbor Hamming-weight-preserving gates on a classical simulator, acting on the alternating product state $|01 \cdots 01\rangle$. The diagonal operator consists of single RZ rotations as well as RZZ interactions between next neighbors, and the time evolution is implemented with a single step of a second-order Trotter formula.

Error mitigation is a crucial technique for leveraging the computational power of noisy quantum devices. Our protocol relies on a noise renormalization method, which has proven effective in prior works [95–98]. The central idea is to assume a depolarizing noise model, estimate its parameters directly on the quantum device, and then invert its effect during post-processing.

We recall that a depolarizing noise model $\mathcal{N}_p[\cdot]$ characterized by a parameter p acts on the expectation value of a Pauli observable σ under a quantum state ρ as

$$\text{Tr}(\sigma \mathcal{N}_p[\rho]) = (1 - p)\text{Tr}(\sigma \rho). \quad (39)$$

Provided the parameter p is known, one can correct the noisy expectation values by dividing them by the factor $(1 - p)$. Here, we can simply estimate this value by replacing the alternating state with the all-zero state $|0 \cdots 0\rangle$, which is invariant under the VQE ansatz. The time evolution is mitigated by performing the first half of the product formula in the forward direction and the second one backward. This method is only effective under depolarizing noise, which is not always an accurate representation of real devices. However, the situation can be improved using Pauli twirling [99]. Readout errors are mitigated by calibrating the device [100], as well as twirling the measurement channel [101].

We provide the results for $L = 10$ qubits, with and without the DB-DOI step, in Table III. State-vector simulations assume ideal, noiseless evolution and infinite sampling, i.e., no shot noise. To parametrize the error mitigation, we used 50 twirls with 1000 shots each, placed on the 156-qubit chip `ibm_fez`. We observe that error mitigation is particularly important when performing the DB-DOI iteration, which

TABLE III. Energy computed on the superconducting quantum chip `ibm_fez` with one DB-DOI step on top of a two-layer warm start. The error bars correspond to a 95% confidence interval on the mean computed via Bayesian data augmentation [97]. The target energy is -14.36 .

	Warm start	1 DB-DOI step
CZ depth	4	22
State vector	-14.01	-14.27
Raw	-13.14 ± 0.01	-10.60 ± 0.01
Mitigated	-13.95 ± 0.01	-14.16 ± 0.04

TABLE IV. Simulation results on Quantinuum’s H1 emulators showing the energy expectation of the open-boundary Heisenberg XXZ model with $L = 10$ and $\Delta = 0.5$ after a two-layer VQE warm start followed by $k = 1$ and $k = 2$ DB-DOI steps. Both the VQE and Hamiltonian simulation are adapted to the platform’s supported native gates. The target energy is of -15.28 .

	Warm start	1 DB-DOI step	2 DB-DOI steps
State vector	-14.16	-14.62	-14.76
H1-1LE	-14.19 ± 0.16	-14.63 ± 0.09	-14.75 ± 0.14
H1-Emulator	-13.93 ± 0.20	-14.59 ± 0.23	-13.65 ± 0.12

significantly increases the depth, as we can outperform the bare VQE.

V. NUMERICAL EMULATION OF DB-DOI FOR QUANTINUUM’S QUANTUM HARDWARE

Our experiment on IBM’s superconducting quantum chip suggests that error mitigation is required for DB-DOI to outperform VQE. However, this requirement may not hold on other quantum platforms, whose distinct noise profiles, coherence times, and native gate fidelities can alter the relative performance of the quantum algorithms. We exemplify this by moving from execution using remotely accessed superconducting qubits to numerical emulation of a trapped-ion platform. Specifically, we use the emulation methods provided as black box by Quantinuum’s Nexus platform. More precisely, our results here are obtained from two types of simulation. The first will be referred to as “H1-1LE” and performs noiseless state-vector simulations that account for sampling (shot) noise. The second approach will be labeled as “H1-Emulator” and accounts for Quantinuum hardware-specific noise model whose parameters are intended to closely mimic the H1 hardware performance. Thus, the data reported in this section are based on numerical simulations. This allows for debugging and optimization of circuits and also for evaluating the performance of our method if executed on quantum hardware. We remark that while these emulators are expected to closely approximate hardware behavior, on the one hand, measurements from a physical processor could exhibit higher noise levels than our simulations; on the other hand, newer hardware models could outperform our numerical predictions.

For instance, Table IV contains the numerical simulation results executed on Quantinuum Nexus, and the model “H1-Emulator” reflects the noise characteristics of Quantinuum’s H1 hardware. Note that we used open-boundary conditions for the execution on IBM quantum hardware in Sec. IV, whereas here we use periodic boundary conditions. Consequently, the target energy is -15.28 compared to the previous -14.36 . We make this choice to present a slightly different example that fits naturally with Quantinuum’s all-to-all qubit connectivity.

We found that one step of DB-DOI is predicted to improve on VQE, despite the expected hardware noise, and no further postprocessing is required. To achieve that, it was important to adapt both the VQE warm start and Hamiltonian simulations in subsequent DB-DOI steps to the native gates supported by the platform. This reduces the physical run-time and hence exposure to decoherence. Additionally, when compiling with

native ZZ gates, any erroneous calibration of the ZZ-gate duration only results in a longer or shorter XXZ Hamiltonian simulation duration during the DB-DOI step; as in both cases they are related to the required cost-function Hamiltonian, the performance impact is effectively due to modifying the DB-DOI step duration s . This is advantageous to the performance of both VQE and DBQA compared to using CZ gates, which require very precise timing of the underlying native ZZ gates to achieve an accurate implementation of Hamiltonian simulation.

In addition, we find that although two DB-DOI steps produce a larger energy drop in the ideal (noise-free) setting, in the presence of realistic device noise, the energy gain is substantially diminished. This is likely because the circuit depth of DB-DOI grew to be large enough to compound decoherence and accumulate gate errors. This is thus an example of the break-even range when using state-of-the-art quantum hardware without quantum error correction.

VI. CONCLUSIONS

Current quantum hardware supports circuits with dozens to hundreds of CZ gates per qubit. Nevertheless, preparing high-fidelity ground-state approximations remains challenging due to the absence of effective compilation methods: VQEs suffer from training limitations in large architectures, while QPE requires fault tolerance. We suggest that the recently introduced DBQA framework [17] offers a promising route to compiling circuits that leverage the capabilities of near-term devices.

Our results indicate that combining VQE with DBQA yields lower energies with shorter training times compared to VQE alone. This hybrid approach may be particularly advantageous on early fault-tolerant architectures, when powerful quantum processors will be remotely accessible. Our implementation [102] within Qibo [103] provides circuit compilations compatible with any QASM-based interface, making experimental demonstrations readily available.

A limitation of our approach is its reliance on sufficiently accurate initializations, which become harder to obtain as system size grows. However, Fig. 3 shows that DBQA markedly improves “*undertrained*” initializations, reducing the energy even when applied for only one or two steps. Moreover, DBQA is agnostic to the initialization method and can be combined with diverse strategies; see Fig. 4 for DB-DOI warm started by HVA on up to 20 qubits, and Appendix F for additional comparisons. Related two-step approaches, such as quantum subspace expansion [104,105], are also natural candidates for early fault-tolerant hardware and stand to benefit from warm starts.

Although rigorous convergence guarantees for DB-DOI require circuit depths that grow exponentially with the number of steps, our numerics demonstrate substantial improvements with only a few steps, provided the initial state has energy below the first excited level. This makes DBQAs particularly suited to early fault-tolerant devices, where shallow circuits are essential but algorithmic guarantees remain valuable. Looking ahead, DBQAs may also serve as an effective initialization for QPE. In this context, quantum dynamic

programming [53] could further reduce circuit depth, enabling DBQAs to act as a bridging tool between near-term and fault-tolerant eras.

Finally, it would be worthwhile to investigate how DBQAs could more broadly enhance variational optimization. An alternative direction to this work would be to use DBQA as a warm start for variational algorithms. Such an approach would connect naturally with a wide range of optimization techniques—Riemannian gradient flows [48,106], natural gradients [107], and quantum imaginary-time evolution [37], among others—opening opportunities to combine their complementary strengths.

ACKNOWLEDGMENTS

We acknowledge the use of IBM Quantum services for this work. We thank K. Bharti, G. Crognaletti, A. Ghanesh, K. Mitarai, A. Sopena, and K. Yamamoto for useful discussions. M.R. and O.K. were supported by the CERN Doctoral Program through the CERN Quantum Technology Initiative. J.S. and N.H.Y.N. were supported by the start-up grant for Nanyang Assistant Professorship of Nanyang Technological University, Singapore. X.L., K.U.G., and M.G. were supported by the Presidential Postdoctoral Fellowship of the Nanyang Technological University, Singapore. J.K. gratefully acknowledges support from Dr. Max Rössler, the Walter Haefner Foundation, and the ETH Zürich Foundation. Z.H. acknowledges support from the Sandoz Family Foundation-Monique de Meuron program for Academic Promotion. This work was further supported by the National Research Foundation Singapore, under its Quantum Engineering Programme 2.0 (National Quantum Computing Hub, NRF2021-QEP2-02-P01). S.T.G. and J.Y.K. acknowledge funding support from A*STAR C230917003. This work was further supported by the National Research Foundation Singapore, under its Quantum Engineering Programme 2.0 (National Quantum Computing Hub, NRF2021-QEP2-02-P01), Quantum Engineering Programme 3.0 Funding Initiative (W24Q3D0002) and Hybrid Quantum-Classical Computing 1.0 Funding Initiative (S24Q7D7001).

The views expressed are those of the authors and do not necessarily reflect the official policy or position of IBM, the IBM Quantum team, or Quantinuum.

DATA AVAILABILITY

All results presented in this work are reproducible using the open-source code at Ref. [102].

APPENDIX A: TACKLING DIFFERENT XXZ SIZES

In this Appendix, we consider one-dimensional XXZ models involving 4–12 qubits. For each case analyzed, we train a four-layer Hamming-weight-preserving ansatz implementing the architecture introduced in Fig. 2 with an Adam optimizer for 300 iterations. The optimization details are the same presented in detail in Appendix C. We repeat the training 10 times for each configuration, initializing the VQE circuit with a different set of angles sampled from the uniform distribution $U(-\pi, \pi)$. We apply three steps of preconditioned DB-DOI

loading each trained VQE at epochs 100, 200, and 300. Each DB-DOI has been optimized through 50 iterations of a CMA-ES [89] algorithm (see Appendix C for DB-DOI optimization details).

The obtained results are shown in Table V and highlight how the DB-DOI approach starts to be impactful on the ground-state preparation with $L \geq 8$ qubits. In fact, the application of DB-DOI after the VQE training allows us to reach a remarkable increase of the fidelity lower bounds even considering only one DB-DOI step. For $L < 8$ qubits, the VQE can be used as standalone because it suffices to obtain the almost exact ground state. Alongside the fidelity lower bounds, we show the depth of the executed circuits computed as the total number of two-qubit gates per qubit. We observe that the depth of the circuits remains approximately constant. This is due to how we compile the DB-DOI in terms of quantum gates. In fact, the contributions due to VQE absorption in the compilation are fixed by the number of layers, and the size of the oracles implementing the Hamiltonian and diagonal evolutions are fixed by the number of time steps and the order of the respective Trotter-Suzuki decomposition. Importantly, these results are obtained assuming each of the VQE circuits composed of four layers. Increasing the size of the problem may require larger VQE circuits, or—in general—more expensive warm-start approximation unitaries. In those cases, the depth of the required DB-DOI circuit will suffer of the scalability inherited from the preconditioning technique.

APPENDIX B: VQE USING HARDWARE-EFFICIENT ANSATZ

In this Appendix, we show how DBQA warm started by a VQE circuit performs when choosing the specific hardware-efficient ansatz targeting the XXZ Hamiltonian. This ansatz is, in principle, more expressive of the Hamming-weight-preserving ansatz of Fig. 2, in the sense that it could cover the whole Hilbert space if the chosen circuit is expressive enough. On one hand, this means the optimal solution to our problem lies in the search space of our algorithm. On the other hand, however, the training is more difficult because no prior knowledge is exploited as it is when choosing the Hamming-weight-preserving ansatz. Practically, we are also exploring regions of the Hilbert space that do not contain the target solution. Comparing the gates composing our chosen HEA circuit with the one introduced in Fig. 2, we note the former presents more parameters but a lower number of CZ. The choice of the ansatz must take into account the hardware availability and calibration accuracy.

We proceed with the same strategy presented in the main text, but considering $N = 5$ trainings per configuration. In Fig. 5, we show an example of DB-DOI execution with a fixed random seed. After an initially rapid cost-function decrease, the VQE training saturates to essentially marginal improvements regardless of the number of layers considered. On the other hand, we find that the DBQA halves the energy residue if applied to very early training epochs (1000, 2000) and essentially reaches the ground state when executed later in the process.

Figure 5 shows an interesting VQE training progression that features two plateaus, the second having an improved cost

TABLE V. Fidelity lower bound $\mathcal{F}$ [see Eq. (C2)] for the XXZ model of Eq. (30) with number of qubits $L \in \{4, 6, 8, 10, 12\}$ and $\Delta = 0.5$. Standalone VQE data are presented alongside DB-DOI results, obtained through compiled DB-DOI circuits. We show the depth of the utilized circuit in the left subcolumn for each experiment setup. The estimates with their uncertainties were calculated using the median and the median absolute deviation of a sample of results obtained by repeating the execution 10 times with different initial conditions.

Qubits	VQE training		1 DB-DOI step		2 DB-DOI steps		3 DB-DOI steps	
	Depth	$\mathcal{F}$	Depth	$\mathcal{F}$	Depth	$\mathcal{F}$	Depth	$\mathcal{F}$
4	16	1	95	1	490	1	2465	1
6	15	1	88	1	457	1	2298	1
8	16	0.947 ± 0.024	95	0.9977 ± 0.0017	490	0.9990 ± 0.0008	2465	0.9995 ± 0.0005
10	16	0.889 ± 0.016	95	0.9928 ± 0.0012	490	0.9972 ± 0.0012	2465	0.9987 ± 0.0005
12	15	0.600 ± 0.091	92	0.919 ± 0.067	473	0.941 ± 0.048	2382	0.956 ± 0.037

function compared to the earlier plateau. We find that even if DB-DOI is initialized in the first plateau it progresses toward the ground-state fixed point in that $k = 1, 2, 3$ iterations of DB-DOI each improve the cost function. When DB-DOI is initialized in the second plateau, then similarly the cost function systematically decreases. In the latter case, the gains are sizable when plotted on a logarithmic axis for the relative difference from the ground-state energy, suggesting rapid convergence.

The lower left plot shows that DBQA initialized in the first training plateau is in the basin of attraction for the ground-state fixed point, where convergence is exponentially fast in the number of steps.

The plot on the right shows the cumulative number of CZ gates required to execute the DB-DOI protocol together with the circuit depth, expressed as the number of CZ gates per qubit. We find that the training cost is dominated by VQE if considering a couple of DBQA steps. On the other hand, the DBQA training optimizes the energy more with fewer queries. For tasks requiring high ground-state preparation fidelity, both methods should be used in sequence as individually each

would necessitate unnecessarily large token expenditures. A more detailed analysis of the performances is presented in Tables VI and VII, where the approximation accuracy is, respectively, quantified in terms of relative difference (RD) or see Eq. (C1) and fidelity lower bound of Eq. (C2).

APPENDIX C: DETAILS ABOUT NUMERICAL SIMULATIONS

In this Appendix, we discuss the choices we made in carrying out the simulations that led to the results shown in this article. After detailing the numerical simulations, we describe the computational cost of the VQE and DB-DOI approaches, which are reported in the tables of this work.

All the simulations have been performed using Qibo [103,108–113], an open-source quantum computing framework widely used to run quantum machine learning algorithms both in simulation [114–120] and on quantum hardware [121–123]. The optimizations have been performed through the Qibo interface, which integrates robust Python




FIG. 5. One example of DB-DOI for XXZ using a hardware-efficient ansatz and obtained by fixing the simulation random seed; the image is intended to provide qualitative information about the impact of DB-DOI. A more robust study of the performance is presented in Table VI. (left) Training of VQE (blue lines) for seven, eight, and nine layers (hues of blue) achieved ground-state energy residue of about 1% within 2500 training epochs. We initialize DBQA with VQE for selected epochs $\in \{1000, 2000, 3000, 4000, 5000\}$, where we apply a DBQA optimized in its parameters with CMA-ES [89]. (right) Token cost estimates of DB-DOI by counting the total number of two-qubit gates required to execute the complete protocol: training the VQE until a target epoch and then optimizing and applying the DBQA.

TABLE VI. (above) Energy approximation ratio for the XXZ model of Eq. (30) with $L = 10$ qubits and $\Delta = 0.5$. The estimates with their uncertainties were calculated using the median and the median absolute deviation of a sample of results obtained by repeating the execution 50 times with different initial conditions. (below) Circuit depth expressed as number of CZ gates per qubit, alongside with cumulative number of CZ gates used to reach $\tilde{E}_0$ (see Appendix C). Warm-start VQE approximations (3000 epochs of training) are presented alongside DB-DOI results, executed considering compiled DB-DOI circuits. Longer VQE training (5000 epochs) is reported in the last column of the table.

Layers	Warm start $1 - \tilde{E}_0/E_0$		1 DB-DOI step $1 - \tilde{E}_0/E_0$		2 DB-DOI steps $1 - \tilde{E}_0/E_0$		Long VQE training $1 - \tilde{E}_0/E_0$	
7	0.023 ± 0.002		0.013 ± 0.0007		0.011 ± 0.0003		0.006 ± 0.002	
8	0.005 ± 0.003		0.0003 ± 0.0002		0.00013 ± 0.00007		0.004 ± 0.002	
9	0.005 ± 0.003		0.0003 ± 0.0002		0.00018 ± 0.00006		0.0030 ± 0.0003	
	Depth	Cumulative cost	Depth	Cumulative cost	Depth	Cumulative cost	Depth	Cumulative cost
7	7	12.18×10^7	50	12.25×10^7	265	12.52×10^7	7	20.03×10^7
8	8	15.84×10^7	55	15.91×10^7	290	16.21×10^7	8	26.4×10^7
9	9	19.98×10^7	60	20.06×10^7	315	20.39×10^7	9	33.3×10^7

based frameworks such as keras [124], tensorflow [125], scipy [126], and pyqma [89].

1. VQE training and computational cost

Training description. We train the VQE using a gradient-based optimization approach. We use the Adam optimizer [91], which has been proven to be one of the most effective optimizers when training big machine-learning models. We set the learning rate to 0.05 after performing a hyperoptimization on a grid of values between 0.1 and 0.001. We keep the default values of the remaining hyperparameters according to the Keras' implementation.

The explored VQE architectures are the Hamming-weight-preserving ansatz introduced in Fig. 2 and a hardware-efficient ansatz composed of ry, rz, and CZ gates. The target Hamiltonians have been chosen as cases of the general Heisenberg Hamiltonian: A first target involves nearest-neighbor interactions and a penalty $\Delta = 0.5$ to the $\hat{Z}$ s, while in a second moment we consider a more complex case lighting up the next-nearest-neighbor interactions. We refer to these two cases as XXZ and J_1 - J_2 , respectively.

Once the VQE ansatz and the target Hamiltonian are specified, we explore its expressibility by varying the number of layers in the VQE ansatz. We trained both ansätze with layers ranging from 3 to 9, and we repeated each training with five different initial configurations. Namely, we randomly sampled the initial values of the angles parametrizing the variational model from the uniform distribution $\mathcal{U}(-\pi, \pi)$.

After selecting the circuit sizes for which the ground-state approximations were most accurate, we conducted a series of more exhaustive simulations to quantify the training error. In particular, we handle the XXZ target training Hamming-weight-preserving ansätze composed of three, four, and five

layers for 2000 epochs, and hardware-efficient ansätze composed of seven, eight, and nine layers for 5000 epochs. We instead tackle the J_1 - J_2 target training Hamming-weight-preserving ansätze composed of three, four, five, and six layers for 2000 epochs.

For each fixed target, ansatz, and number of layers, we repeat the training 50 times, with a different set of initial parameters sampled from the uniform distribution $\mathcal{U}(-\pi, \pi)$. Each of these training instances corresponds to a ground-state energy approximation $\tilde{E}_0$, whose quality is quantified by computing the relative difference (RD) with the target (known or numerically computed) ground-state energy E_0 :

$$\text{RD} = 1 - \frac{\tilde{E}_0}{E_0}. \quad (\text{C1})$$

Alternatively, we can use the ground-state energy approximations $\tilde{E}_0$ to compute the following fidelity lower bound [69]:

$$\mathcal{F} = 1 - \frac{\tilde{E}_0 - E_0}{E_1 - E_0}, \quad (\text{C2})$$

where E_0 and E_1 are the true ground state and first excited state, respectively.

Once all the 50 RD or $\mathcal{F}$ values are collected, a final estimate is computed through the median value of the list, $\text{median}(\mathbf{x})$, where we call $\mathbf{x}$ the list of RD/ $\mathcal{F}$ values for simplicity. The estimation of the uncertainty is instead computed using the median absolute deviation:

$$\text{MAD} = 1.4826 \times \text{median}(|x_i - \text{median}(\mathbf{x})|). \quad (\text{C3})$$

We choose the median and the median absolute deviation as estimators of the variable and its uncertainty to be more robust to outliers, which could be expected when applying the DB-DOI algorithm. Since we make use of a global cost

TABLE VII. Fidelity lower bound $\mathcal{F}$ [see Eq. (C2)] for the XXZ model of Eq. (30) with $L = 10$ qubits and $\Delta = 0.5$ extending results of Table VI.

Layers	Warm start	1 DB-DOI step	2 DB-DOI steps	Long VQE training
7	0.67 ± 0.03	0.82 ± 0.05	0.85 ± 0.05	0.92 ± 0.03
8	0.93 ± 0.04	0.995 ± 0.002	0.998 ± 0.001	0.94 ± 0.03
9	0.93 ± 0.02	0.996 ± 0.001	0.998 ± 0.001	0.96 ± 0.004

function in this work, it can happen that, if the initial approximation provided by the VQE is not close enough to the ground state, the final effect of the DBQA does not correspond to a further reduction of the energy. Moreover, since the optimization cost of the DBQA is particularly intense in the case of the noncompiled DBI, in some rare cases it can happen that the optimizers are not able to find an optimal configuration of the parameters when the optimization process is limited in time.

Computational cost. To evaluate the computational cost of the VQE training, we need to take into account the total number of two-qubit gates composing the circuit architecture and the number of times an expectation value has to be computed to evaluate predictions and gradients during the training. When using a gradient-based approach on a quantum device, the gradients must be calculated using parameter shift rules [90,127]. Since we parametrize the quantum circuit through rotational gates, it is well known that the partial derivative of our cost function with respect to a circuit's parameter θ can be calculated using two expectation values. Considering the choice of Adam optimization, the entire gradient of the cost function has to be computed at each optimization iteration. Finally, the total amount of two-qubit gates (we take CZ as reference) can be computed as

$$N_{\text{CZ}}^{\text{VQE}}(e) = k \text{pen}_{\text{CZ}}^{\text{VQE}}, \quad (\text{C4})$$

where we indicate with e the number of Adam iterations, p the number of parameters of the circuit, and $n_{\text{CZ}}^{\text{VQE}}$ the number of CZ gates that compose the circuit according to the chosen ansatz. The constant k refers to the number of expectation values required to execute the parameter shift rule. In the case of the hardware-efficient ansatz, $k = 2$, and in the case of the Hamming-weight-preserving ansatz, $k = 4$, since each RBS gate is decomposed into two rotations depending on the same angle θ [83]. According to the same decomposition, the number of CZ gates composing the Hamming-weight-preserving ansatz circuit can be calculated as twice the number of RBS gates.

2. DB-DOI procedure and computational cost

DB-DOI procedure. To fully exploit the algorithm, it is necessary to prepare an approximation of the ground state. In this work, we make use of a VQE, but any ground-state preparation algorithm can be used. This first approximation is then used to precondition the target Hamiltonian, which is then rotated according to Eq. (14). The following key step consists in compiling the DBI circuit and this can be done following an equivalent procedure to the one presented in Appendix D. In the following discussion of the computational cost of the algorithm, we will take into account the same VQE trainings exposed in the previous section, which are used as preconditioning ground-state approximations in the DB-DOI process. We evaluate Eq. (C1) after applying the DB-DOI to a given ground-state approximation provided by VQE stopped at target epoch e . Following the same procedure of Appendix C 1, we collect all the RD values and provide the final estimations as the median and median absolute deviation of the obtained results.

Computational cost. Both the initial state approximation and the DB-DOI compilation into a circuit present a computational cost and have to be cumulatively taken into account to evaluate the whole computational expense of the algorithm. The total number of two-qubit gates required by the DB-DOI process involves then a first contribution due to the VQE cost described in Eq. (C4). We have to consider then the cost of executing the new circuit, in which the VQE unitary is repeated exponentially as DBI iterations increase, and there is an additional cost due to the Hamiltonian decomposition as explained in Appendix D. Denoting $n_{\text{CZ}}^{\text{DBQA}}$ the number of CZ gates composing the circuit already involving both the DB-DOI compilation and the recursive VQE call, we finally must consider the cost of optimizing the DB-DOI's parameters. In fact, the DB-DOI circuit can be parametrized both in the step duration s and in the diagonal operators $\hat{D}_k$. In this work, we parametrize $\hat{D}_k$ as a classical Ising model with NN interactions:

$$\hat{D}_k(B^{(k)}, J^{(k)}) = \sum_{i=0}^N (B_i^{(k)} \hat{Z}_i + J_i^{(k)} \hat{Z}_{i+1} \hat{Z}_i), \quad (\text{C5})$$

where N is the number of sites in the chain, the parametrization is implemented through the coefficients $\{\alpha_i\}_{i=0}^N$ and $\{\beta_i\}_{i=0}^N$, and the superscript k highlights that the optimization procedure is repeated for each DBQA step. We use Scipy's Powell and CMA-ES optimizers to find an optimal DBI configuration, and this optimization has to be considered as an additional cost in the overall count. To evaluate this final contribution to the total number of two-qubit gates, we multiply $n_{\text{CZ}}^{\text{DBQA}}$ with the total number of cost-function evaluations n_{fval} executed by the optimizers. The total number of two-qubit gates required to finalize the DB-DOI ground-state approximation can be finally obtained as

$$N_{\text{CZ}}^{\text{DBQA}} = N_{\text{CZ}}^{\text{VQE}}(e) + n_{\text{fval}} n_{\text{CZ}}^{\text{DBQA}}. \quad (\text{C6})$$

APPENDIX D: COMPILING OF XXZ EVOLUTION

We next provide a discussion on how to explicitly express the diagonalization DBQA as above into an explicit circuit. For the XXZ model, we write the Hamiltonian as

$$\hat{H}_0 = \sum_{a=1}^L \hat{H}^{(a)}, \quad (\text{D1})$$

where each summand addresses only two qubits

$$\hat{H}^{(a)} = \hat{X}_a \hat{X}_{a+1} + \hat{Y}_a \hat{Y}_{a+1} + \hat{Z}_a \hat{Z}_{a+1} \quad (\text{D2})$$

for $a < L$ and

$$\hat{H}^{(L)} = \hat{X}_L \hat{X}_1 + \hat{Y}_L \hat{Y}_1 + \hat{Z}_L \hat{Z}_1 \quad (\text{D3})$$

because of periodic boundary conditions. We next use M steps of the linear Trotter-Suzuki decomposition

$$e^{-it\hat{H}_0} = \left(\prod_{a=1}^L e^{-i\frac{t}{M}\hat{H}^{(a)}} \right)^M + O(M^{-1}), \quad (\text{D4})$$

which means that, if we decompose $e^{-it/M\hat{H}^{(a)}}$ into CNOT and single-qubit rotations, we get a circuit approximation with accuracy $O(t^2/M)$. This is done by noticing that $e^{-it/M\hat{H}^{(a)}}$

is a unitary acting on two qubits, i.e., a and $a + 1$, and any two-qubit unitary can be decomposed into a circuit made of single-qubit rotations and three CNOT gates [128].

We next discuss generalizations. The Hamiltonian in Eq. (D1) can be arbitrary and only the last step has to be modified: Then, the terms $\hat{H}^{(a)}$ can be acting on more than two qubits and the evolutions they generate $e^{-it/M\hat{H}^{(a)}}$ can be compiled into CNOT and single-qubit rotations using the quantum Shannon decomposition [129] that implements a unitary on K qubits using $O(4^K)$ CNOTs that is, in general, an optimal scaling.

Additionally, we can use a higher-order Trotter-Suzuki decomposition. We assume again that we have a two-qubit Hamiltonian on a line of an even number of qubits L so that $\hat{H}_0^{(o)} = \sum_{a=1}^{L/2} \hat{H}^{(2a-1)}$ and $\hat{H}_0^{(e)}$ contains only commuting terms. Then, we write

$$e^{-it\hat{H}_0} = \left(e^{-it/(2M)\hat{H}_0^{(o)}} e^{-it/M\hat{H}_0^{(e)}} e^{-it/(2M)\hat{H}_0^{(o)}} \right)^M + O(M^{-2}), \quad (\text{D5})$$

which gives an improved error scaling. From commutativity, we have the exact equality

$$e^{-it/(2M)\hat{H}_0^{(o)}} = \prod_{a=1}^{L/2} e^{-it/M\hat{H}^{(2a+1)}}. \quad (\text{D6})$$

Each of the factors above can again be decomposed into CNOTs and single-qubit rotations by standard methods [128,129].

1. Other 2-local models

The addition of the magnetic field Hamiltonian can be compiled in the above formalism by setting

$$\hat{H}^{(a)} = \hat{X}_a \hat{X}_{a+1} + \hat{Y}_a \hat{Y}_{a+1} + \hat{Z}_a \hat{Z}_{a+1} + B_a \hat{Z}_a \quad (\text{D7})$$

for $a < L$ and for periodic boundary conditions, we set

$$\hat{H}^{(L)} = \hat{X}_L \hat{X}_1 + \hat{Y}_L \hat{Y}_1 + \hat{Z}_L \hat{Z}_1 + B_L \hat{Z}_L \quad (\text{D8})$$

or for open-boundary conditions $\hat{H}^{(L)} = 0$ so that

$$\hat{H}_0 = \hat{H}_{\text{XXZ}} + \hat{H}(B) = \sum_{a=1}^L \hat{H}^{(a)}. \quad (\text{D9})$$

For the transverse-longitudinal field Ising model, we set

$$\hat{H}^{(a)} = \hat{X}_a \hat{X}_{a+1} + B_a \hat{Z}_a + C_a \hat{X}_a \quad (\text{D10})$$

and boundary terms similar to those above. This would give as above a compiling solution with three CNOT gates.

2. Special-purpose compiling for the transverse-field Ising model

We next consider the special transverse-field Ising model

$$\hat{H}^{(a)} = \hat{X}_a \hat{X}_{a+1} + B_a \hat{Z}_a \quad (\text{D11})$$

obtained from the above by setting $C_a = 0$. We use that $\hat{X}_a \hat{X}_b = \text{CNOT}(a, b) \hat{X}_a \text{CNOT}(a, b)$ and $\hat{Z}_a = \text{CNOT}(a, b) \hat{Z}_a \text{CNOT}(a, b)$ to write

$$\text{CNOT}(a, b) \hat{H}^{(a)} \text{CNOT}(a, b) = \hat{X}_a + B_a \hat{Z}_a, \quad (\text{D12})$$

which means that we can get for any t

$$e^{-it\hat{H}^{(a)}} = \text{CNOT}(a, b) e^{-it(\hat{X}_a + B_a \hat{Z}_a)} \text{CNOT}(a, b). \quad (\text{D13})$$

In other words, the transverse-field Ising model evolutions can be compiled using two, not three, CNOT gates per interaction term.

3. Compiling for the classical Ising model

In the numerical calculations, we use parametrizations for the diagonal evolutions that have a low quantum compiling. More specifically, we define the classical Ising model

$$\hat{H}(B, J) = \sum_{a=1}^L (B_a \hat{Z}_a + J_{a,a+1} \hat{Z}_a \hat{Z}_{a+1}), \quad (\text{D14})$$

where as above $\hat{Z}_{L+1} = \hat{Z}_1$. For $\hat{H}(B, J)$, we can compile the diagonal evolution using two CNOT gates, which is more efficient than using the general method of compiling two-qubit unitaries from Ref. [128]. We use that $\hat{Z}_a \hat{Z}_b = \text{CNOT}(a, b) \hat{Z}_b \text{CNOT}(a, b)$ and by unitarity of the CNOT gate,

$$e^{-it\hat{Z}_a \hat{Z}_b} = \text{CNOT}(a, b) e^{-it\hat{Z}_b} \text{CNOT}(a, b). \quad (\text{D15})$$

The model is commuting so $e^{-it\hat{H}(B, J=0)}$ consists of independent single-qubit rotations.

APPENDIX E: DB-DOI RESULTS CONSIDERING THE J_1 - J_2 MODEL

In the main text, we presented results for the SU(2) symmetric XXZ model, where we set $\Delta = 0.5$ as a penalty to the $\hat{Z}$ interactions. Here, we extend the analysis to a more general model:

$$\hat{H}_{J_1-J_2} = J_1 \hat{H}_{\text{XXZ}} + J_2 \sum_{i=1}^L (\hat{X}_i \hat{X}_{i+2} + \hat{Y}_i \hat{Y}_{i+2} + \hat{Z}_i \hat{Z}_{i+2}), \quad (\text{E1})$$

where we light up next-nearest-neighbor interactions. In particular, we consider $J_1 = 1$ and $J_2 = 0.2$, which corresponds to a regime such that $J_2/J_1 = 0.2$, and has been chosen knowing the target system presents a Berezinskii-Kosterlitz-Thouless transition at $J_2/J_1 = 0.24116$ [130,131]. Also, in this case, we use the Hamming-weight-preserving ansatz presented in Fig. 2 because the same symmetries are respected. We follow the same procedure as the one presented in the main text for XXZ, but in this case, we execute the DBQA in single commutator mode only. We postpone to future works the compilation of the model into a quantum circuit, likewise we did in Appendix D. The obtained results are presented in Table VIII.

APPENDIX F: DETAILS ON RELATION TO OTHER METHODS

In this Appendix, we provide a supplementary discussion of the relation to other methods for a low-energy state preparation of a quantum many-body system. It is meant to gather contributions to the literature that may be helpful for using DBQAs in the future when tackling large systems will be required and conjunction with other methods will become particularly relevant.

TABLE VIII. Relative difference between approximated energy $\tilde{E}_0$ and the target ground-state value E_0 for the J_1 - J_2 model of Eq. (E1) with $L = 10$ qubits, $J_1 = 1$, and $J_2 = 0.2$. The estimates with their uncertainties were calculated using the median and the median absolute deviation of a sample of results obtained by repeating the execution 50 times with different initial conditions. Warm-start VQE approximations (500 epochs of training) are presented alongside DB-DOI results, executed considering DBI unitaries. Longer VQE trainings (2000 epochs) are reported in the last column of the table.

Layers	Warm start	1 DBI step	2 DBI steps	3 DBI steps	Long VQE training
3	0.033 ± 0.005	0.020 ± 0.005	0.016 ± 0.005	0.014 ± 0.004	0.026 ± 0.006
4	0.017 ± 0.007	0.006 ± 0.005	0.003 ± 0.003	0.002 ± 0.002	0.010 ± 0.004
5	0.011 ± 0.007	0.002 ± 0.002	0.0008 ± 0.0008	0.0004 ± 0.0004	0.005 ± 0.003
6	0.009 ± 0.006	0.002 ± 0.001	0.0007 ± 0.0007	0.0004 ± 0.0004	0.005 ± 0.003

Exploring other VQE warm starts. There exist many other VQE ideas, and we refer to reviews for specific discussion of their detailed performance [8–10]. Overall, it is clear that purely variational methods can be limited by training obstructions (e.g., swamps of local minima or barren plateaus). So, the accuracy in energy estimation is expected to saturate before reaching the global minimum. In all variational cases, one can interface the unitary $\hat{U}_\theta$ s with DBQA via the rotation of the input generator as done in the main text.

Classical learning of Riemannian gradients. We note that DB-DOI is conceptually distinct from Ref. [106], which also uses methods involving exponentials of commutators, framed there in the language of Riemannian flows. The difference is that Ref. [106] focuses on computing the gradients of generic fixed parametrized quantum circuits, while we work with the full manifold rather than a submanifold. This eases implementations of state updates in gradient directions and unlike in Ref. [106] measurements are not required in DB-DOI. We refer to Ref. [48] for an early review Riemannian gradients and Riemannian gradients as an underlying principle for performing quantum control and quantum computing tasks. References [21,26] studied mathematically iterations involving exponentials of commutators under the name Lie-bracket recursions (while we refer to them as double-bracket iterations for consistency with double-bracket flows referring to the continuous flows) and proved their convergence.

Gradient-based optimization is often used for training VQE, and it is natural to consider these methods for training DBQA as well. An example of such one-dimensional optimization is to use a greedy strategy to select DBI step durations s_k as cost-function minimizers for the respective double-bracket rotation. This strategy can be more general, e.g., gradient descent in the space of magnetic fields B_i or Ising couplings $J_{i,j}$ to find the best $\hat{D}_k$ operators in each step, starting from an initial guess. We have also tried nongradient-based optimizers (see Appendix C). See Refs. [106,132–134] for a discussion of Riemannian gradients projected onto particular subspaces, e.g., those spanned by randomly chosen Pauli operators.

Analytic warm starts. An interesting avenue of research is to replace VQE unitaries $\hat{U}_\theta$ by Hartree-Fock warm starts; see Ref. [32] for a discussion of quantum compiling using Givens rotations and experimental data. In this case, again, the question would be whether Hartree-Fock interfaced with DBQA could achieve better ground-state approximation. Further generalization to other ansätze in quantum chemistry can

be interfaced with DBQA as long as explicit compiling is available.

Variationally learned QITE warm starts. Apart from direct minimization of energy, there exist methods that aim to implement certain transformations, and these transformations have energy-lowering properties. Quantum imaginary-time evolution is an iteration where in each step one seeks an approximation to the imaginary-time update $|\psi_{k+1}\rangle = e^{-\tau\hat{H}_0}|\psi_k\rangle / \|e^{-\tau\hat{H}_0}|\psi_k\rangle\|$ through a unitary $\hat{U}_k$ such that $\hat{U}_k|\psi_k\rangle \approx |\psi_{k+1}\rangle$. Here, the optimal parametrization of $\hat{U}_k$ is not known so variational methods have been used [35,37,38]. QITE circumvents trainability issues by trying to implement steps of an iteration that is bound to decrease the energy, but it is not clear how to find the correct unitary for longer durations where the cost-function resolution is limiting. Instead, DBQA is coherent in that it can be applied directly and only relies on measurements to optimize its performance but does not rely on classical supervision for its parametrization. For this reason, it is likely that the variationally learned QITE approximation circuits as in Refs. [35,37,38] should come before DBQA rather than after. However, if it is possible to easily parametrize QITE at late stages of energy optimization, then warm-starting DBQA by QITE from Refs. [35,37,38] and then following again by training based on a QITE training proposed in Refs. [35,37,38] could be useful too.

Relation of DBQA to fault-tolerant quantum algorithms. So far, we discussed methods that have been suggested without the use of auxiliary qubits. If one allows for that, then methods based on quantum phase estimation become key. Here again, interfacing can be possible. Most likely DBQA should precede QPE rather than following QPE by DBQA because QPE can be expected to be more costly than DBQA, so it is natural to chain methods to increase resource cost, e.g., start with VQE, then perform a few steps of DBQA before switching to QPE. QPE methods can be systematically formulated in the formalism of qubitization. Note that QITE can also be implemented in this way. In all cases, the upper bound on the ground-state energy found via DB-DOI (or composition of other methods) could also be complemented with a lower bound of the ground-state energy via the dual-VQE proposal [135].

Relation of DBQA and QAOA. Finally, let us comment broadly on DBQA concerning the so-called *quantum approximate optimization algorithm* (QAOA) [8,136]. Let us begin with Eq. (17) and notice the resemblance to a one-layer QAOA, in that in $\hat{U}_1$ we have a layer of single-qubit

gates (specifically, diagonal Z rotations) followed by the evolution under the input Hamiltonian and then again single-qubit rotations. For the first step, this is similar to QAOA that interlaces single-qubit rotation layers with evolutions under the input Hamiltonian. The biggest difference is that in DBQA, the two single-qubit layers are related in that they are inverses of each other. Thus, the guidance of the DBI equation through the reduced group commutator approximation suggests a more constrained QAOA ansatz. This constraint facilitates the monotonicity relation [17], which allows us to analytically understand why diagonalization ensues. Similar relations for unconstrained QAOA are not known to us, but Ref. [17] conjectured that this might hint at understanding the functioning of optimized QAOAs. Aside from these coincidental similarities, the methods are different in several key aspects. In Ref. [136], the QAOA ansatz involving evolutions with the native Hamiltonian for trapped ions was used to prepare an approximation to the ground state of a different model. In contrast, DBQA is not restricted to target only the ground state but can be applied to target any other eigenstate of the given model. The distinction from QAOA becomes more evident when applying DBQA beyond the first step. For example, at second order the unitary in Eq. (23) reveals an additional feature: $\hat{U}_2$ involves not only forward evolutions but also backward. Indeed, while it may be challenging to implement DBQA beyond the first step on analog quantum simulators, it is this nested forward and backward evolution structure that ensures the convergence of DBQA. We are not aware of convergence guarantees of the DBQA type for QAOA.

1. Overview of currently available circuit depth

Following the overview above, most methods for preparing ground states on existing quantum hardware suffer from one of two severe problems. Variational methods are limited by the resolution of the cost function, which is at the root of their operation. This limits the circuit depth that is meaningful in practice to about a dozen entangling layers. QPE-based methods are the opposite in that to become meaningful, they need a large circuit depth.

As a specific example, we consider Ref. [137], which compared the performance of QPE for physical and logical qubits. The algorithmic performance in Ref. [137] was obtained by using $N \approx 920$ CZ gates for $L = 8$ qubits, and the two-qubit gate fidelity was $p_e \approx 2 \times 10^{-3}$.

We can use these experimental results to get ballpark figures for currently available depths. Upcoming devices are set to reach $p'_e \approx 5 \times 10^{-4}$ [138], so assuming each CZ gate fails independently of the others, we can estimate the number N' of gates that can be used meaningfully by solving heuristically

$$(1 - p_e)^N = (1 - p'_e)^{N'}. \quad (\text{F1})$$

This equation gives 16% success probability, which in the three-nines fidelity regime p'_e would appear for $N' \approx 3690$ CZ gates. From these estimates, we see that there exists quantum hardware that operates meaningfully for circuit depths of about 100 CZ gates per qubit, and this is set to even larger values in the very near future. This means that one and two DB-DOI steps are already feasible for implementation on existing noisy intermediate-scale quantum devices.

-
- [1] Y. Dong, L. Lin, and Y. Tong, Ground-state preparation and energy estimation on early fault-tolerant quantum computers via quantum eigenvalue transformation of unitary matrices, *PRX Quantum* **3**, 040305 (2022).
 - [2] L. Lin and Y. Tong, Heisenberg-limited ground-state energy estimation for early fault-tolerant quantum computers, *PRX Quantum* **3**, 010318 (2022).
 - [3] Y. Ge, J. Tura, and J. I. Cirac, Faster ground state preparation and high-precision ground energy estimation with fewer qubits, *J. Math. Phys.* **60**, 022202 (2019).
 - [4] O. Kiss, U. Azad, B. Requena, A. Roggero, D. Wakeham, and J. M. Arrazola, Early fault-tolerant quantum algorithms in practice: Application to ground-state energy estimation, *Quantum* **9**, 1682 (2025).
 - [5] H. R. Grimsley, S. E. Economou, E. Barnes, and N. J. Mayhall, An adaptive variational algorithm for exact molecular simulations on a quantum computer, *Nat. Commun.* **10**, 3007 (2019).
 - [6] D. W. Berry, Y. Tong, T. Khattar, A. White, T. I. Kim, G. H. Low, S. Boixo, Z. Ding, L. Lin, S. Lee, G. K.-L. Chan, R. Babbush, and N. C. Rubin, Rapid initial-state preparation for the quantum simulation of strongly correlated molecules, *PRX Quantum* **6**, 020327 (2025).
 - [7] L. Lin and Y. Tong, Near-optimal ground state preparation, *Quantum* **4**, 372 (2020).
 - [8] K. Bharti *et al.*, Noisy intermediate-scale quantum algorithms, *Rev. Mod. Phys.* **94**, 015004 (2022).
 - [9] M. Cerezo, A. Arrasmith, R. Babbush, S. C. Benjamin, S. Endo, K. Fujii, J. R. McClean, K. Mitarai, X. Yuan, L. Cincio, *et al.*, Variational quantum algorithms, *Nat. Rev. Phys.* **3**, 625 (2021).
 - [10] J. Tilly, H. Chen, S. Cao, D. Picozzi, K. Setia, Y. Li, E. Grant, L. Wossnig, I. Rungger, G. H. Booth, *et al.*, The variational quantum eigensolver: A review of methods and best practices, *Phys. Rep.* **986**, 1 (2022).
 - [11] E. R. Anschuetz and B. T. Kiani, Beyond barren plateaus: Quantum variational algorithms are swamped with traps, *Nat. Commun.* **13**, 7760 (2022).
 - [12] J. R. McClean, S. Boixo, V. N. Smelyanskiy, R. Babbush, and H. Neven, Barren plateaus in quantum neural network training landscapes, *Nat. Commun.* **9**, 4812 (2018).
 - [13] L. Bittel and M. Kliesch, Training variational quantum algorithms is NP-Hard, *Phys. Rev. Lett.* **127**, 120502 (2021).
 - [14] D. Stilck Franca and R. Garcia-Patron, Limitations of optimization algorithms on noisy quantum devices, *Nat. Phys.* **17**, 1221 (2021).
 - [15] M. Larocca, S. Thanasilp, S. Wang, K. Sharma, J. Biamonte, P. J. Coles, L. Cincio, J. R. McClean, Z. Holmes, and M. Cerezo, A review of barren plateaus in variational quantum computing, *Nat. Rev. Phys.* **7**, 174 (2024).
 - [16] M. Cerezo, M. Larocca, D. García-Martín, N. L. Diaz, P. Braccia, E. Fontana, M. S. Rudolph, P. Bermejo, A. Ijaz, S. Thanasilp, *et al.*, Does provable absence of barren plateaus

- imply classical simulability? Or, why we need to rethink variational quantum computing, *Nat. Commun.* **16**, 7907 (2023).
- [17] M. Gluza, Double-bracket quantum algorithms for diagonalization, *Quantum* **8**, 1316 (2024).
 - [18] M. Gluza, J. Son, B. H. Tiang, R. Zander, R. Seidel, Y. Suzuki, Z. Holmes, and N. H. Y. Ng, Double-bracket quantum algorithms for quantum imaginary-time evolution, *Phys. Rev. Lett.* **136**, 020601 (2026).
 - [19] R. Zander, R. Seidel, L. Xiaoyue, and M. Gluza, Role of Riemannian geometry in double-bracket quantum imaginary-time evolution, in *Geometric Science of Information*, edited by F. Nielsen and F. Barbaresco (Springer Nature Switzerland, Cham, 2026), pp. 105–114.
 - [20] Y. Suzuki, B. H. Tiang, J. Son, N. H. Y. Ng, Z. Holmes, and M. Gluza, Double-bracket algorithm for quantum signal processing without post-selection, *Quantum* **9**, 1954 (2025).
 - [21] U. Helmke and J. B. Moore, *Optimization and Dynamical Systems* (Springer, London, 1994).
 - [22] Y. Suzuki, M. Gluza, J. Son, B. H. Tiang, N. H. Ng, and Z. Holmes, Grover's algorithm is an approximation of imaginary-time evolution, [arXiv:2507.15065](https://arxiv.org/abs/2507.15065).
 - [23] A. Katabarwa, K. Gratsea, A. Caesura, and P. D. Johnson, Early fault-tolerant quantum computing, *PRX Quantum* **5**, 020101 (2024).
 - [24] J. Preskill, Quantum computing in the NISQ era and beyond, *Quantum* **2**, 79 (2018).
 - [25] J. Montanez-Barrera, K. Michielsen, and D. E. B. Neira, Evaluating the performance of quantum processing units at large width and depth, [arXiv:2502.06471](https://arxiv.org/abs/2502.06471).
 - [26] J. B. Moore, R. E. Mahony, and U. Helmke, Numerical gradient algorithms for eigenvalue and singular value calculations, *SIAM J. Matrix Anal. Appl.* **15**, 881 (1994).
 - [27] S. T. Smith, *Geometric Optimization Methods for Adaptive Filtering* (Harvard University, Cambridge, MA, 1993).
 - [28] A. Peruzzo, J. McClean, P. Shadbolt, M.-H. Yung, X.-Q. Zhou, P. J. Love, A. Aspuru-Guzik, and J. L. O'Brien, A variational eigenvalue solver on a photonic quantum processor, *Nat. Commun.* **5**, 4213 (2014).
 - [29] J. L. Bosse and A. Montanaro, Probing ground-state properties of the kagome antiferromagnetic Heisenberg model using the variational quantum eigensolver, *Phys. Rev. B* **105**, 094409 (2022).
 - [30] J. Kattemölle and J. van Wezel, Variational quantum eigensolver for the Heisenberg antiferromagnet on the kagome lattice, *Phys. Rev. B* **106**, 214429 (2022).
 - [31] D. Wecker, M. B. Hastings, and M. Troyer, Progress towards practical quantum variational algorithms, *Phys. Rev. A* **92**, 042303 (2015).
 - [32] F. Arute, K. Arya, R. Babbush, D. Bacon, J. C. Bardin, R. Barends, S. Boixo, M. Broughton, B. B. Buckley, D. A. Buell, *et al.*, Hartree-Fock on a superconducting qubit quantum computer, *Science* **369**, 1084 (2020).
 - [33] S. R. White, Density matrix formulation for quantum renormalization groups, *Phys. Rev. Lett.* **69**, 2863 (1992).
 - [34] K. G. Wilson, The renormalization group: Critical phenomena and the Kondo problem, *Rev. Mod. Phys.* **47**, 773 (1975).
 - [35] M. Motta, C. Sun, Adrian T. K. Tan, M. J. O'Rourke, E. Ye, A. J. Minnich, F. G. S. L. Brandão, and G. K.-L. Chan, Determining eigenstates and thermal states on a quantum computer using quantum imaginary time evolution, *Nat. Phys.* **16**, 205 (2020).
 - [36] H. Nishi, T. Kosugi, and Y.-i. Matsushita, Implementation of quantum imaginary-time evolution method on NISQ devices by introducing nonlocal approximation, *npj Quantum Inf.* **7**, 85 (2021).
 - [37] S. McArdle, T. Jones, S. Endo, Y. Li, S. C. Benjamin, and X. Yuan, Variational ansatz-based quantum simulation of imaginary time evolution, *npj Quantum Inf.* **5**, 75 (2019).
 - [38] N. Gomes, A. Mukherjee, F. Zhang, T. Iadecola, C.-Z. Wang, K.-M. Ho, P. P. Orth, and Y.-X. Yao, Adaptive variational quantum imaginary time evolution approach for ground state preparation, *Adv. Quantum Technol.* **4**, 2100114 (2021).
 - [39] S. Fomichev, K. Hejazi, M. S. Zini, M. Kiser, J. Fraxanet, P. A. M. Casares, A. Delgado, J. Huh, A.-C. Voigt, J. E. Mueller, and J. M. Arrazola, Initial state preparation for quantum chemistry on quantum computers, *PRX Quantum* **5**, 040339 (2024).
 - [40] F. Verstraete, M. Wolf, and J. Ignacio Cirac, Quantum computation and quantum-state engineering driven by dissipation, *Nat. Phys.* **5**, 633 (2009).
 - [41] M.-A. Filip and N. Fitzpatrick, Beyond asymptotic reasoning: The practicalities of a quantum ground state projector based on the Wall–Chebyshev expansion, *Quantum Sci. Technol.* **11**, 015027 (2026).
 - [42] F. Wegner, Flow-equations for Hamiltonians, *Ann. Phys.* **506**, 77 (1994).
 - [43] S. D. Glazek and K. G. Wilson, Renormalization of Hamiltonians, *Phys. Rev. D* **48**, 5863 (1993).
 - [44] F. Wegner, Flow equations and normal ordering: A survey, *J. Phys. A: Math. Gen.* **39**, 8221 (2006).
 - [45] S. Kehrein, *The Flow Equation Approach to Many-Particle Systems*, Springer Tracts in Modern Physics, Vol. 217 (Springer, Berlin, 2006).
 - [46] R. Brockett, Dynamical systems that sort lists, diagonalize matrices, and solve linear programming problems, *Linear Algebra Appl.* **146**, 79 (1991).
 - [47] A. Anshu, D. Gosset, K. J. Morenz Korol, and M. Soleimanifar, Improved approximation algorithms for bounded-degree local Hamiltonians, *Phys. Rev. Lett.* **127**, 250502 (2021).
 - [48] T. Schulte-Herbrüggen, S. J. Glaser, G. Dirr, and U. Helmke, Gradient flows for optimization in quantum information and quantum dynamics: Foundations and applications, *Rev. Math. Phys.* **22**, 597 (2010).
 - [49] B. Zindorf and S. Bose, Efficient implementation of multi-controlled quantum gates, *Phys. Rev. Appl.* **24**, 044030 (2025).
 - [50] S. Balaucă and A. Arusoae, Efficient constructions for simulating multi controlled quantum gates, in *International Conference on Computational Science* (Springer, Cham, Switzerland, 2022), pp. 179–194.
 - [51] B. Zindorf and S. Bose, Multi-controlled quantum gates in linear nearest neighbor, [arXiv:2506.00695](https://arxiv.org/abs/2506.00695).
 - [52] X. Wang, A. Sorensen, and K. Molmer, Multibit gates for quantum computing, *Phys. Rev. Lett.* **86**, 3907 (2001).
 - [53] J. Son, M. Gluza, R. Takagi, and N. H. Y. Ng, Quantum dynamic programming *Phys. Rev. Lett.* **134**, 180602 (2024).
 - [54] Y.-A. Chen, A. M. Childs, M. Hafezi, Z. Jiang, H. Kim, and Y. Xu, Efficient product formulas for commutators and

- applications to quantum simulation, *Phys. Rev. Res.* **4**, 013191 (2022).
- [55] A. M. Childs and N. Wiebe, Product formulas for exponentials of commutators, *J. Math. Phys.* **54**, 062202 (2013).
- [56] G. Kuperberg, Breaking the cubic barrier in the Solovay–Kitaev algorithm, [arXiv:2306.13158](#).
- [57] A. Elkasapy and A. Thom, On the length of the shortest non-trivial element in the derived and the lower central series, *J. Group Theory* **18**, 793 (2015).
- [58] A. Elkasapy, A new construction for the shortest non-trivial element in the lower central series, [arXiv:1610.09725](#).
- [59] J. Peetz and P. Narang, Hamiltonian simulation via stochastic Zassenhaus expansions, [arXiv:2501.13922](#).
- [60] A. M. Childs and Y. Su, Nearly optimal lattice simulation by product formulas, *Phys. Rev. Lett.* **123**, 050503 (2019).
- [61] E. Campbell, Random compiler for fast Hamiltonian simulation, *Phys. Rev. Lett.* **123**, 070503 (2019).
- [62] L. Xiaoyue, M. Robbiati, A. Pasquale, E. Pedicillo, A. Wright, S. Carrazza, and M. Gluza, Strategies for optimizing double-bracket quantum algorithms, [arXiv:2408.07431](#).
- [63] F. Levkovich-Maslyuk, The Bethe ansatz, *J. Phys. A: Math. Theor.* **49**, 323004 (2016).
- [64] A. Sopena, M. H. Gordon, D. García-Martín, G. Sierra, and E. López, Algebraic Bethe circuits, *Quantum* **6**, 796 (2022).
- [65] R. Ruiz, A. Sopena, M. H. Gordon, G. Sierra, and E. López, The Bethe ansatz as a quantum circuit, *Quantum* **8**, 1356 (2024).
- [66] V. Murg, V. E. Korepin, and F. Verstraete, Algebraic Bethe ansatz and tensor networks, *Phys. Rev. B* **86**, 045125 (2012).
- [67] R. I. Nepomechie, Bethe ansatz on a quantum computer? [arXiv:2010.01609](#).
- [68] T. Fogarty, S. Deffner, T. Busch, and S. Campbell, Orthogonality catastrophe as a consequence of the quantum speed limit, *Phys. Rev. Lett.* **124**, 110601 (2020).
- [69] M. Cramer, M. B. Plenio, S. T. Flammia, R. Somma, D. Gross, S. D. Bartlett, O. Landon-Cardinal, D. Poulin, and Y.-K. Liu, Efficient quantum state tomography, *Nat. Commun.* **1**, 149 (2010).
- [70] Z. Holmes, K. Sharma, M. Cerezo, and P. J. Coles, Connecting ansatz expressibility to gradient magnitudes and barren plateaus, *PRX Quantum* **3**, 010313 (2022).
- [71] P. Nataf and F. Mila, Exact diagonalization of Heisenberg $SU(N)$ models, *Phys. Rev. Lett.* **113**, 127204 (2014).
- [72] M. Raissi, P. Perdikaris, and G. Karniadakis, Physics-informed neural networks: A deep learning framework for solving forward and inverse problems involving nonlinear partial differential equations, *J. Comput. Phys.* **378**, 686 (2019).
- [73] G.-L. R. Anselmetti, D. Wierichs, C. Gogolin and R. M. Parrish, Local, expressive, quantum-number-preserving VQE ansätze for fermionic systems, *New J. Phys.* **23**, 113010 (2021).
- [74] J. M. Arrazola, O. Di Matteo, N. Quesada, S. Jahangiri, A. Delgado, and N. Killoran, Universal quantum circuits for quantum chemistry, *Quantum* **6**, 742 (2022).
- [75] S. Johri, S. Debnath, A. Mocherla, A. Singk, A. Prakash, J. Kim, and I. Kerenidis, Nearest centroid classification on a trapped ion quantum computer, *npj Quantum Inf.* **7**, 122 (2021).
- [76] I. Kerenidis and A. Prakash, Quantum machine learning with subspace states, [arXiv:2202.00054](#).
- [77] N. Jain, J. Landman, N. Mathur, and I. Kerenidis, Quantum Fourier networks for solving parametric PDEs, *Quantum Sci. Technol.* **9**, 035026 (2024).
- [78] M. Ragone, P. Braccia, Q. T. Nguyen, L. Schatzki, P. J. Coles, F. Sauvage, M. Larocca, and M. Cerezo, Representation theory for geometric quantum machine learning, [arXiv:2210.07980](#).
- [79] L. Monbroussou, J. Landman, A. B. Grilo, R. Kukla, and E. Kashefi, Trainability and expressivity of Hamming-weight preserving quantum circuits for machine learning, *Quantum* **9**, 1745 (2023).
- [80] C. Tüysüz, S. Y. Chang, M. Demidik, K. Jansen, S. Vallecorsa, and M. Grossi, Symmetry breaking in geometric quantum machine learning in the presence of noise *PRX Quantum* **5**, 030314 (2024).
- [81] D. Raveh and R. I. Nepomechie, Estimating Bethe roots with VQE, *J. Phys. A: Math. Theor.* **57**, 355303 (2024).
- [82] R. M. S. Farias, T. O. Maciel, G. Camilo, R. Lin, S. Ramos-Calderer, and L. Aolita, Quantum encoder for fixed Hamming-weight subspaces, *Phys. Rev. Appl.* **23**, 044014 (2025).
- [83] E. A. Cherrat, I. Kerenidis, N. Mathur, J. Landman, M. Strahm, and Y. Y. Li, Quantum vision transformers, *Quantum* **8**, 1265 (2024).
- [84] G. Crognaletti, G. D. Bartolomeo, M. Vischi, and L. L. Viteritti, Equivariant variational quantum eigensolver to detect phase transitions through energy level crossings, *Quantum Sci. Technol.* **10**, 015048 (2025).
- [85] O. Kiss, M. Grossi, P. Lougovski, F. Sanchez, S. Vallecorsa, and T. Papenbrock, Quantum computing of the ${}^6\text{Li}$ nucleus via ordered unitary coupled clusters, *Phys. Rev. C* **106**, 034325 (2022).
- [86] C. Gu, M. Heinz, O. Kiss, and T. Papenbrock, Toward scalable quantum computations of atomic nuclei, *Phys. Rev. C* **113**, 034321 (2026).
- [87] M. Grossi, O. Kiss, F. De Luca, C. Zollo, I. Gremese, and A. Mandarino, Finite-size criticality in fully connected spin models on superconducting quantum hardware, *Phys. Rev. E* **107**, 024113 (2023).
- [88] A. Kandala, A. Mezzacapo, K. Temme, M. Takita, M. Brink, J. M. Chow, and J. M. Gambetta, Hardware-efficient variational quantum eigensolver for small molecules and quantum magnets, *Nature (London)* **549**, 242 (2017).
- [89] N. Hansen, The CMA evolution strategy: A tutorial, [arXiv:1604.00772](#).
- [90] M. Schuld, V. Bergholm, C. Gogolin, J. Izaac, and N. Killoran, Evaluating analytic gradients on quantum hardware, *Phys. Rev. A* **99**, 032331 (2019).
- [91] D. P. Kingma and J. Ba, in *International Conference on Learning Representations (ICLR)*, (2015).
- [92] C.-Y. Park and N. Killoran, Hamiltonian variational ansatz without barren plateaus, *Quantum* **8**, 1239 (2024).
- [93] R. Wiersema, C. Zhou, Y. de Sereville, J. F. Carrasquilla, Y. B. Kim, and H. Yuen, Exploring entanglement and optimization within the Hamiltonian variational ansatz, *PRX Quantum* **1**, 020319 (2020).
- [94] M. Larocca, P. Czarnik, K. Sharma, G. Muraleedharan, P. J. Coles, and M. Cerezo, Diagnosing barren plateaus with tools from quantum optimal control, *Quantum* **6**, 824 (2022).
- [95] O. Kiss, D. Teplitskiy, M. Grossi, and A. Mandarino, Statistics of topological defects across a phase transition in a

- superconducting quantum processor, *Quantum Sci. Technol.* **10**, 035037 (2025).
- [96] M. Urbanek, B. Nachman, V. R. Pascuzzi, A. He, C. W. Bauer, and W. A. de Jong, Mitigating depolarizing noise on quantum computers with noise-estimation circuits, *Phys. Rev. Lett.* **127**, 270502 (2021).
- [97] O. Kiss, M. Grossi, and A. Roggero, Quantum error mitigation for Fourier moment computation, *Phys. Rev. D* **111**, 034504 (2025).
- [98] O. Kiss, I. Tavernelli, F. Tacchino, D. Lacroix, and A. Roggero, Neutrino thermalization via randomization on a quantum processor, [arXiv:2510.24841](https://arxiv.org/abs/2510.24841).
- [99] Z. Cai and S. Benjamin, Constructing smaller Pauli twirling sets for arbitrary error channels, *Sci. Rep.* **9**, 1128 (2019).
- [100] P. D. Nation, H. Kang, N. Sundaresan, and J. M. Gambetta, Scalable mitigation of measurement errors on quantum computers, *PRX Quantum* **2**, 040326 (2021).
- [101] E. van den Berg, Z. K. Mineev, and K. Temme, Model-free readout-error mitigation for quantum expectation values, *Phys. Rev. A* **105**, 032620 (2022).
- [102] Double-bracket quantum algorithms in Qibo, GitHub, 2025, <https://github.com/qiboteam/qibo-dbqa>.
- [103] S. Efthymiou, S. Ramos-Calderer, C. Bravo-Prieto, A. Pérez-Salinas, D. García-Martín, A. García-Saiz, J. I. Latorre, and S. Carrazza, Qibo: A framework for quantum simulation with hardware acceleration, *Quantum Sci. Technol.* **7**, 015018 (2022).
- [104] J. Robledo-Moreno, M. Motta, H. Haas, A. Javadi-Abhari, P. Jurcevic, W. Kirby, S. Martiel, K. Sharma, S. Sharma, T. Shirakawa, I. Sitdikov, R.-Y. Sun, K. J. Sung, M. Takita, M. C. Tran, S. Yunoki, and A. Mezzacapo, Chemistry beyond the scale of exact diagonalization on a quantum-centric supercomputer, *Sci. Adv.* **11**, eadu9991 (2025).
- [105] N. Yoshioka, M. Amico, W. Kirby, P. Jurcevic, A. Dutt, B. Fuller, S. Garion, H. Haas, I. Hamamura, A. Ivrii, R. Majumdar, Z. Mineev, M. Motta, B. Pokharel, P. Rivero, K. Sharma, C. J. Wood, A. Javadi-Abhari, and A. Mezzacapo, Krylov diagonalization of large many-body Hamiltonians on a quantum processor, *Nat. Commun.* **16**, 5014 (2025).
- [106] R. Wiersema and N. Killoran, Optimizing quantum circuits with Riemannian gradient flow, *Phys. Rev. A* **107**, 062421 (2023).
- [107] J. Stokes, J. Isaac, N. Killoran, and G. Carleo, Quantum natural gradient, *Quantum* **4**, 269 (2020).
- [108] S. Carrazza, S. Efthymiou, M. Lazzarin, and A. Pasquale, An open-source modular framework for quantum computing, *J. Phys.: Conf. Ser.* **2438**, 012148 (2023).
- [109] S. Efthymiou, M. Lazzarin, A. Pasquale, and S. Carrazza, Quantum simulation with just-in-time compilation, *Quantum* **6**, 814 (2022).
- [110] S. Efthymiou *et al.*, Qibolab: An open-source hybrid quantum operating system, *Quantum* **8**, 1247 (2024).
- [111] A. Pasquale, A. Papaluca, R. M. S. Farias, M. Robbiati, E. Pedicillo, and S. Carrazza, Beyond full statevector simulation with Qibo, [arXiv:2408.00384](https://arxiv.org/abs/2408.00384).
- [112] E. Pedicillo, A. Candido, S. Efthymiou, H. Sargsyan, Y. P. Tan, J. Cereijo, J. Y. Khoo, A. Pasquale, M. Robbiati, and S. Carrazza, An open-source framework for quantum hardware control, [arXiv:2407.21737](https://arxiv.org/abs/2407.21737).
- [113] A. Pasquale, S. Efthymiou, S. Ramos-Calderer, J. Wilkens, I. Roth, and S. Carrazza, Towards an open-source framework to perform quantum calibration and characterization, [arXiv:2303.10397](https://arxiv.org/abs/2303.10397).
- [114] A. Pérez-Salinas, J. Cruz-Martínez, A. A. Alhajri, and S. Carrazza, Determining the proton content with a quantum computer, *Phys. Rev. D* **103**, 034027 (2021).
- [115] C. Bravo-Prieto, J. Baglio, M. Cè, A. Francis, D. M. Grabowska, and S. Carrazza, Style-based quantum generative adversarial networks for Monte Carlo events, *Quantum* **6**, 777 (2022).
- [116] M. Robbiati, J. M. Cruz-Martínez, and S. Carrazza, Determining probability density functions with adiabatic quantum computing, *Quantum Mach. Intell.* **7**, 5 (2025).
- [117] J. M. Cruz-Martínez, M. Robbiati, and S. Carrazza, Multi-variable integration with a variational quantum circuit, *Quantum Sci. Technol.* **9**, 035053 (2024).
- [118] S. Bordoni, D. Stanev, T. Santantonio, and S. Giagu, Long-lived particles anomaly detection with parametrized quantum circuits, *Particles* **6**, 297 (2023).
- [119] V. Belis, K. A. Woźniak, E. Puljak, P. Barkoutsos, G. Dissertori, M. Grossi, M. Pierini, F. Reiter, I. Tavernelli, and S. Vallecorsa, Quantum anomaly detection in the latent space of proton collision events at the LHC, *Commun. Phys.* **7**, 334 (2024).
- [120] J. J. Martínez de Lejarza, L. Cieri, M. Grossi, S. Vallecorsa, and G. Rodrigo, Loop Feynman integration on a quantum computer, *Phys. Rev. D* **110**, 074031 (2024).
- [121] M. Robbiati, S. Efthymiou, A. Pasquale, and S. Carrazza, A quantum analytical Adam descent through parameter shift rule using Qibo, [arXiv:2210.10787](https://arxiv.org/abs/2210.10787).
- [122] M. Robbiati, A. Sopena, A. Papaluca, and S. Carrazza, Real-time error mitigation for variational optimization on quantum hardware, [arXiv:2311.05680](https://arxiv.org/abs/2311.05680).
- [123] A. D'Elia *et al.*, Characterization of a transmon qubit in a 3D cavity for quantum machine learning and photon counting, *Appl. Sci.* **14**, 1478 (2024).
- [124] F. Chollet *et al.*, Keras (2015), <https://keras.io>.
- [125] M. Abadi *et al.*, TensorFlow: Large-scale machine learning on heterogeneous systems (2015), software available from <http://tensorflow.org>.
- [126] P. Virtanen *et al.*, SciPy 1.0: Fundamental algorithms for scientific computing in Python, *Nat. Methods* **17**, 261 (2020).
- [127] K. Mitarai, M. Negoro, M. Kitagawa, and K. Fujii, Quantum circuit learning, *Phys. Rev. A* **98**, 032309 (2018).
- [128] F. Vatan and C. Williams, Optimal quantum circuits for general two-qubit gates, *Phys. Rev. A* **69**, 032315 (2004).
- [129] V. V. Shende, S. S. Bullock, and I. L. Markov, Synthesis of quantum logic circuits, in *Proceedings of the 2005 Asia and South Pacific Design Automation Conference* (ACM Digital, New York, 2005), pp. 272–275.
- [130] J. V. José, *40 Years of Berezinskii-Kosterlitz-Thouless Theory* (World Scientific, Singapore, 2013).
- [131] S. R. White and I. Affleck, Dimerization and incommensurate spiral spin correlations in the zigzag spin chain: Analogies to the Kondo lattice, *Phys. Rev. B* **54**, 9862 (1996).

- [132] S. Bordoni *et al.*, Quantum noise modeling through reinforcement learning, [Quantum Sci. Technol.](#) **11**, 015005 (2026).
- [133] N. A. McMahon, M. Pervez, and C. Arenz, Equating quantum imaginary time evolution, Riemannian gradient flows, and stochastic implementations, [Phys. Rev. Res.](#) **8**, 023024 (2026).
- [134] Z. Lai, H. Nie, J. Wu, and D. An, Quantum circuit design from a retraction-based Riemannian optimization framework, [arXiv:2602.20605](#).
- [135] H. Westerheim, J. Chen, Z. Holmes, I. Luo, T. Nuradha, D. Patel, S. Rethinasamy, K. Wang, and M. M. Wilde, Dual variational quantum eigensolver: A quantum algorithm to lower bound the ground-state energy, [Phys. Rev. A](#) **113**, 032443 (2026).
- [136] C. Kokail, C. Maier, R. van Bijnen, T. Brydges, M. K. Joshi, P. Jurcevic, C. A. Muschik, P. Silvi, R. Blatt, C. F. Roos, *et al.*, Self-verifying variational quantum simulation of lattice models, [Nature \(London\)](#) **569**, 355 (2019).
- [137] K. Yamamoto, S. Duffield, Y. Kikuchi, and D. M. Ramo, Demonstrating Bayesian quantum phase estimation with quantum error detection, [Phys. Rev. Res.](#) **6**, 013221 (2024).
- [138] K. Yamamoto (private communication).