

Imaginary time spectral transforms for excited-state preparation

D. A. Millar^{1,2,*}, L. W. Anderson¹, E. Altamura^{3,4}, O. Wallis³, M. E. Sahin³, J. Crain^{5,6} and S. J. Thomson^{1,7,†}¹IBM Research, Hursley, Winchester SO21 2JN, United Kingdom²School of Physics and Astronomy, University of Southampton, Southampton SO17 1BJ, United Kingdom³The Hartree Centre, STFC, Sci-Tech Daresbury, Warrington WA4 4AD, United Kingdom⁴Yusuf Hamied Department of Chemistry, University of Cambridge, Lensfield Road, Cambridge CB2 1EW, United Kingdom⁵IBM Research Europe, The Hartree Centre, Sci-Tech Daresbury, Warrington WA4 4AD, United Kingdom⁶Clarendon Laboratory, University of Oxford, Oxford OX1 3PU, United Kingdom⁷SUPA, School of Physics and Astronomy, University of Edinburgh, Peter Guthrie Tait Road, Edinburgh EH9 3FD, United Kingdom

(Received 8 August 2025; accepted 3 May 2026; published 3 June 2026)

Excited states of many-body quantum systems play a key role in a wide range of physical and chemical phenomena. Despite this, there is a notable lack of scalable algorithms capable of preparing highly excited eigenstates. To address this challenge, we introduce a general approach that directly targets eigenstates near a chosen energy, applicable to both classical and quantum simulation frameworks. Our approach combines the shift-invert mechanism with imaginary time evolution, enabling the construction of excited states of large many-body quantum systems. We demonstrate the technique classically by computing midspectrum eigenstates of disordered spin chains with up to $L = 128$ sites. Based on this, we propose a hybrid scheme compatible with near-term quantum hardware.

DOI: [10.1103/v4rc-gmcx](https://doi.org/10.1103/v4rc-gmcx)

Introduction. The excited-state manifold of quantum systems is important for a variety of physical and chemical properties. In condensed matter physics, excited states govern key phenomena such as information transport, critical behavior, and thermodynamics in various applications, including molecular magnets, optically addressable spins, quantum dots, and trapped ions [1–4]. Knowledge of the excited-state structure can even allow construction of thermal Gibbs states at low temperatures [5]. In chemistry, molecular excitations can alter reactivities, redox properties, and proton transfer barriers [6,7]. They can even activate reaction pathways and isomerization transitions that are inaccessible in the ground state [8], with potentially transformative consequences for drug discovery and design, as well as affect bonding, geometry, chemical properties, and spectroscopic signatures [9,10]. Similarly, in electronic structure, the excited state of a molecular species often exhibits distinct chemical properties compared to its ground state, influencing bonding, geometry, and spectroscopic signatures [11–13]. Broad classes of materials underpinning optoelectronics and quantum devices, such as those in light-emitting diodes and photovoltaic cells, all rely on harnessing excited-state attributes to confer functionality [14]. Excited states offer a vastly expanded molecular design space, playing a pivotal and growing role in various

applications that rely upon light-matter interactions or charge injection [15].

Despite the importance of excited states, a full understanding of their structure has remained elusive. Unlike ground states, for which many efficient and widely used algorithms exist [including quantum Monte Carlo [16,17] or the density matrix renormalization group (DMRG) [18,19]], there are few ways to compute excited states outside of full exact diagonalization. Constructing low-lying excited states iteratively from the ground state is often feasible, but constructing highly excited states in this manner poses a significant challenge, often requiring the computation of hundreds or thousands of unwanted states before obtaining the desired one. The ability to directly compute highly excited eigenstates—without requiring the explicit construction of all lower-energy states—is highly desirable for many fundamental problems, including those related to eigenstate thermalization and many-body localization [20–22]. Classical tensor network techniques exist for directly computing excited states of large systems in various limits, such as strong disorder [23–25] or weakly entangled states such as quantum many-body scars [26]. Despite impressive results, these algorithms are ultimately hindered by the inability of classical computers to efficiently capture highly entangled states. There is a clear need for a robust algorithm that can be deployed on both classical and quantum computers, allowing us to overcome these limitations and study highly entangled excited states—This is the central challenge we address in this work.

Model. We demonstrate our technique using the disordered Heisenberg model, which is given by

$$\mathcal{H} = J \sum_i \mathbf{S}_i \cdot \mathbf{S}_{i+1} + \sum_i h_i S_i^z, \quad (1)$$

*Contact author: declan.millar@ibm.com†Contact author: steven.thomson@ed.ac.uk




FIG. 1. Schematic of the algorithm. The method combines shift-invert and imaginary time evolution: The shift-invert process maps eigenstates of an initial Hamiltonian $\mathcal{H}$ near a target energy δ to the extremal eigenstates of an effective Hamiltonian $\mathcal{H}_{\text{eff}}$, which are then determined via imaginary time evolution. To avoid explicit inversion of $(\mathcal{H} - \delta)$, we derive a self-consistency condition for an infinitesimal time step $d\tau$, which we solve variationally. We evolve the state iteratively until we reach a target variance. To provide a “warm start” to the algorithm, we initialize it using an approximation to the ground state of $(\mathcal{H} - \delta)^2$.

where $h_i \in [-W, W]$ is the on-site disorder taken from a uniform distribution, and we set $J = 1$ as our unit of energy. We restrict our analysis to $W > 0$, as in the limit of $W \rightarrow 0$, the energy spectrum is highly degenerate. In the presence of degeneracies, our method is likely to converge to a superposition of degenerate states. Other methods (including DMRG) have similar limitations.

Method. Our algorithm, sketched in Fig. 1, combines two core techniques: spectral transforms and imaginary time evolution. Spectral transforms involve modifying the Hamiltonian such that an eigenstate at some target energy δ becomes an extremal eigenstate of a new effective Hamiltonian $\mathcal{H}_{\text{eff}}$. We may then use any reliable ground-state search technique to find the ground state of $\pm \mathcal{H}_{\text{eff}}$. One example is *spectral folding* $\mathcal{H} \rightarrow \mathcal{H}_{\text{eff}} = (\mathcal{H} - \delta)^2$, which has been used in the context of quantum chemistry [27]. Throughout, we abbreviate constant terms like $\mathbb{1}\delta \rightarrow \delta$ for clarity. This transform typically results in a large density of states in the vicinity of the target state, which can be challenging for variational algorithms (results shown in Ref. [28]). This problem can be avoided by instead using the *shift-invert technique* [29]:

$$\mathcal{H} \rightarrow \mathcal{H}_{\text{eff}} = (\mathcal{H} - \delta)^{-1}. \quad (2)$$

The ground state of $\mathcal{H}_{\text{eff}}$ will now be the closest eigenstate of $\mathcal{H}$ with energy less than δ , while the maximally excited state of $\mathcal{H}_{\text{eff}}$ will be the closest state with energy greater than δ . Note that if δ is an exact eigenvalue, the shifted Hamiltonian is singular. This is typically not a problem in practice, as we are extremely unlikely to pick an exact eigenvalue by accident. We can find the highest excited state of $\mathcal{H}_{\text{eff}}$ by solving for the ground state of $-\mathcal{H}_{\text{eff}}$. Since the shift-invert transform corresponds to a resolvent of the Hamiltonian, the procedure may also be viewed as an energy-selective spectral transformation.

The key obstacle in applying the shift-invert method is that direct inversion of the Hamiltonian is costly, with naive methods scaling like $\mathcal{O}(N_H^3)$, where N_H is the dimension of the Hilbert space, which is itself typically exponentially large in the system size. Moreover, even if $\mathcal{H}_{\text{eff}}$ can be computed, it is unlikely to be sparse, making it complicated to directly implement in frameworks such as tensor networks or on near-

term quantum hardware. Sophisticated techniques have been developed to avoid having to invert the Hamiltonian, often involving solving an equivalent implicit problem. These are, however, often specific to particular classical implementations [25,29] and are not convenient for direct implementation on quantum hardware. Here, we avoid explicit inversion of the Hamiltonian in a way that does not depend on any particular implementation framework, either quantum or classical, by combining the shift-invert procedure with *imaginary time evolution*.

Imaginary time evolution is a well-established technique with guarantees for obtaining ground states of quantum systems [30]. It involves preparing the system in an arbitrary initial state, presumed to be far from the ground state of the chosen Hamiltonian—but with nonzero overlap—and evolving in imaginary time τ . This evolution is generated by an operator of the form $e^{-\mathcal{H}\tau}$, where $\mathcal{H}$ is the microscopic Hamiltonian and τ represents imaginary time. The ultimate goal is to prepare the state $|\psi(\tau)\rangle / \| |\psi(\tau)\rangle \|_2$, where $|\psi(\tau)\rangle = e^{-\mathcal{H}\tau} |\psi(0)\rangle$, which in the limit of $\tau \rightarrow \infty$ will be the zero-temperature ground state of $\mathcal{H}$. Our crucial modification to this technique is that we shall perform the imaginary time evolution with respect to the shift-inverted Hamiltonian $\mathcal{H}_{\text{eff}}$, allowing us to use imaginary time evolution to target excited states.

The (un-normalized) time-evolved state is given by

$$|\psi(\tau)\rangle = e^{-\mathcal{H}_{\text{eff}}\tau} |\psi(0)\rangle = e^{-(\mathcal{H}-\delta)^{-1}\tau} |\psi(0)\rangle. \quad (3)$$

We start by making an infinitesimal imaginary time step of size $d\tau$ and expand the evolution operator to obtain

$$|\psi(d\tau + \tau)\rangle \approx (1 - (\mathcal{H} - \delta)^{-1}d\tau)|\psi(\tau)\rangle. \quad (4)$$

We then multiply both sides by $\mathcal{H}_{\text{shift}} = (\mathcal{H} - \delta)$ to obtain

$$\mathcal{H}_{\text{shift}}|\psi(d\tau + \tau)\rangle = (\mathcal{H}_{\text{shift}} - d\tau)|\psi(\tau)\rangle. \quad (5)$$

This gives us a self-consistency condition that can be solved variationally to evolve the state in (imaginary) time $|\psi(\tau)\rangle \rightarrow |\psi(\tau + d\tau)\rangle$. At each time step, we evolve the state by minimizing the Hilbert-Schmidt distance. Other cost functions can be used, some of which are simpler to measure on a quantum

computer; however, we found the Hilbert-Schmidt distance to be significantly more numerically stable than all alternatives we tested:

$$D = \|\mathcal{H}_{\text{shift}} |\psi(\tau + d\tau)\rangle - (\mathcal{H}_{\text{shift}} - d\tau) |\psi(\tau)\rangle\|_2. \quad (6)$$

We have thus traded the problematic matrix inverse for a variational update step that can be performed classically or on near-term quantum hardware [31,32]. The iterative nature of the process ensures that as long as $d\tau$ is small, the convergence properties of the algorithm remain favorable [33], while also inheriting the performance and accuracy guarantees of conventional imaginary time evolution.

Evaluating the cost function in Eq. (6) on a quantum computer requires estimating $\langle\psi(\tau + d\tau)|\mathcal{H}_{\text{shift}}^2|\psi(\tau + d\tau)\rangle$, $\langle\psi(\tau)|(\mathcal{H}_{\text{shift}} - d\tau)^2|\psi(\tau)\rangle$, and $\text{Re}[\langle\psi(\tau + d\tau)|\mathcal{H}_{\text{shift}}(\mathcal{H}_{\text{shift}} - d\tau)|\psi(\tau)\rangle]$. The first two terms are expectation values of positive semidefinite operators and can be estimated using standard measurement routines. The cross term can be written as

$$(\mathcal{H} - \delta)(\mathcal{H} - \delta - d\tau) = \sum_i \alpha_i P_i, \quad (7)$$

where P_i are Pauli strings. This reduces the overlap term to a sum of quantities of the form $\text{Re}[\langle\psi(\tau + d\tau)|P_i|\psi(\tau)\rangle]$, each of which can be measured using a modified Hadamard test [28]. For local Hamiltonians like Eq. (1), this expansion generally contains $\mathcal{O}(L^2)$ Pauli strings, although mutually commuting terms can be grouped to reduce the number of measurement bases. Consequently, for local Hamiltonians the number of grouped observables scales polynomially with system size, as does the sampling cost per variational update. Estimating each grouped observable to additive precision ϵ with failure probability η requires $\mathcal{O}(\log(1/\eta)/\epsilon^2)$ shots. This can be further reduced to $\mathcal{O}(\log(1/\eta)/\epsilon)$ if amplitude amplification is used, albeit at the cost of deeper circuits [34,35]. In this context, the circuit depth would be controlled by the ansatz chosen to represent the state, and would not grow with the total imaginary time τ as in conventional (real-time) Trotter evolution. Instead, increasing τ corresponds to performing additional variational updates of a fixed circuit structure, more akin to the McLachlan variational approach to time evolution. Further details regarding measurement protocols, compilation cost, and noise considerations are discussed in Ref. [28].

We now have all of the essential ingredients for our algorithm (Fig. 1). (1) First, we use standard techniques to obtain a loose approximation to the ground state of the folded Hamiltonian $(\mathcal{H} - \delta)^2$. This could be obtained on quantum hardware using, for example, the variational quantum eigensolver (VQE) [36,37]. Here, we compute it classically using DMRG (typically with initial bond dimension $\chi_0 \leq 4$). The state we obtain will in general not be an exact eigenstate, but will serve as a good initial choice for $|\psi(\tau = 0)\rangle$. This procedure ensures that our initial state has nonzero overlap with eigenstates near the target energy, mitigating against the “orthogonality catastrophe” that would cause randomly chosen states to have a vanishingly small overlap with the target state. (2) We then take $|\psi(\tau = 0)\rangle$ —which could be expressed as a shallow quantum circuit [38–41]—and use it as the input state for the shift-inverted imaginary time evolution. (3) We evolve the system in time by repeatedly applying Eq. (6).

As $\tau \rightarrow \infty$, we asymptotically approach an eigenstate with energy close to δ . The precise final state will depend on the overlap of the initial state with states in the vicinity of the target energy. Closeness to the desired energy can be verified by computing $E(\tau) = \langle\psi(\tau)|\mathcal{H}|\psi(\tau)\rangle$ and comparing with the target energy. We verify closeness to an eigenstate by computing the variance $\sigma(\tau) = \langle\psi(\tau)|(\mathcal{H} - E(\tau))^2|\psi(\tau)\rangle$, which vanishes for an exact eigenstate.

Results. We present a proof-of-concept implementation of the algorithm using a matrix product state (MPS) representation, based on `ITensors.jl` [42]. Details of the implementation are given in Ref. [28]. The use of matrix product states is not essential to the algorithm: It can equally be implemented directly on state vectors [28], directly on quantum circuits, or directly on any other desired framework. While MPS methods are not efficient for highly entangled states, there generically exist states in the bulk of the spectrum with moderate bond dimension that we expect to capture with this approach [43]. To ensure accuracy, we update the size of the time step $d\tau$ such that the Taylor expansion $\exp(-\mathcal{H}_{\text{eff}} d\tau) |\psi(\tau)\rangle \approx (1 - \mathcal{H}_{\text{eff}} d\tau + \mathcal{O}(d\tau^2)) |\psi(\tau)\rangle$ is valid. We do this by computing $E(\tau)$ and choosing $d\tau_{\min} \leq d\tau \ll 2(E(\tau) - \delta)$, imposing the lower limit $d\tau_{\min} = 10^{-3}$ to prevent slowdown when $E(\tau) \approx \delta$.

First, we demonstrate the results of the method for a small, exactly solvable system of size $L = 12$, for a variety of disorder strengths. We evolve the initial state until the fidelity with any exact eigenstate $\mathcal{F} = \max\{|\langle\psi(\tau)|\phi_i\rangle|^2 : i = 1, \dots, 2^L\}$ is at least $\mathcal{F} = 0.999$ or the variance decreases below $\sigma^* = 10^{-6}$. We also compute the von Neumann entanglement entropy $S = -\text{Tr}[\rho_A \log \rho_A] = -\text{Tr}[\rho_B \log \rho_B]$, where subsystems A and B are formed by a bipartition at a given bond, and we average S over all bonds. We set a target energy of $\delta = 0$, which is typically a midspectrum state, representing a challenging area of the spectrum where we can expect high entanglement and a high density of states. The results are shown in Fig. 2, where the overline indicates the average over disorder realizations. These results enable us to (1) demonstrate that the method is capable of recovering the exact solution to high accuracy and (2) show that the midspectrum states we obtain are always more highly entangled than the corresponding ground states (obtained with DMRG), giving us confidence that we are finding representative midspectrum states rather than exceptionally weakly entangled states. We see that the algorithm converges to high accuracy regardless of the disorder strength. As the disorder strength decreases, the entanglement in the midspectrum states grows much faster than in the corresponding ground states. The more highly entangled the target state is, the larger the bond dimension we require, and the more computationally costly it is to solve Eq. (6) for each time step due to the rapidly increasing number of parameters. Ultimately, for highly entangled states in large systems, the bond dimension growth will become prohibitive on classical hardware, eventually necessitating the use of quantum computers to construct such states using this algorithm. Even classically, it may be advantageous to represent states as tailored quantum circuits, which can have far fewer parameters than their MPS representation [44]. Interestingly, we see no clear indication of the many-body localization (MBL) transition close to $W/J \approx 4$ [20,21,45,46], although




FIG. 2. Results for system size $L = 12$ for disorder strengths $W \in [1, 10]$, averaged over $N_s \in [15, 30]$ disorder realizations for each point, depending on the disorder strength. Overlines denote the disorder average, and error bars indicate the standard deviation. (a) The deviation from unit fidelity ($1 - \bar{F}$) of midspectrum states. (b) The final variance $\bar{\sigma}$. (c) The final bond dimension $\bar{\chi}$. (d) Comparison of the bond-averaged von Neumann entanglement entropy of the midspectrum states (red) and the corresponding ground states (blue). Solid lines in panels (b)–(d) are guides to the eye.

the entanglement entropy appears to flatten off above this point, and convergence becomes significantly slower below it. It would be interesting to perform a systematic study of large systems to investigate whether or not a clear signature of an MBL phase transition can be found [22,47–50].

We now show that in the regime of moderate-to-strong disorder where the entanglement growth is inhibited, we can classically construct low-variance states of large systems. We show representative results for both $L = 64$ with $W/J = 6$ and $L = 128$ with $W/J = 10$. We use a demanding cutoff of $\sigma^* < 1 \times 10^{-5}$, which is much lower than the average variance required to reach an eigenstate for $L = 12$ at $W/J = 6$ (see Fig. 2). When comparing different system sizes, it may be fairer to scale the variance by system size [51]; however, we opt not to do so here in order to present our results as transparently as possible. The results are shown in Fig. 3. Using DMRG to extract the extremal eigenvalues of the Hamiltonian, we find that with respect to the target energy $\delta = 0$, the final states shown have a relative energy error of $\Delta E = (E(\tau \rightarrow \infty) - \delta)/(E_{\max} - E_{\min}) \approx 10^{-4}$ and $\Delta E \approx 10^{-5}$, respectively. In this regime, we would expect a very high density of states with level spacings that may be smaller than can be resolved by machine precision. Starting in a low bond dimension MPS partially mitigates against this, as we do not expect this very sparse initial state will have significant overlap with many contiguous eigenstates within a given energy window. Further details are given in the Supplemental Material [28], along with additional results for system sizes $L \in [48, 64]$ with $W/J = 6$ and $L = 128$ with $W/J \in [6, 10]$. Our strategy of starting in an initial MPS with a small bond dimension and allowing $\chi(\tau)$ to grow gradually




FIG. 3. Imaginary time evolution of (a) the energy $E(\tau)$, (b) the variance $\sigma(\tau)$, and (c) the maximum bond dimension. Columns show system size $L = 64$ with disorder strength $W/J = 6$ (left) and $L = 128$ with $W/J = 10$ (right). The dashed red line in panel (a) shows the target energy, $\delta = 0$, and the dashed gray line in panel (b) shows the variance threshold, $\sigma^* = 1 \times 10^{-5}$.

may take more time steps than directly starting with an MPS of large bond dimension, but it is typically significantly faster in terms of elapsed real time. This also suggests a potential hybrid quantum/classical approach, where the initial evolution is performed classically until the bond dimension begins to grow too large. The low bond dimension MPS can then be translated into a shallow quantum circuit [38–41], and the final portion of the evolution performed on a quantum computer.

Discussion and conclusion. In this work, we have described an algorithm for the preparation of excited eigenstates (and low-variance states) of many-body quantum systems. We have shown results obtained numerically on classical hardware and proposed a hybrid implementation compatible with near-term quantum computers. We have showcased the method using a common condensed matter model; however, the technique itself is Hamiltonian-agnostic, requiring no prior knowledge of the problem structure, and will work for a wide variety of many-body quantum systems, including applications in quantum chemistry. Nor does it depend on the simulation method used: The core algorithm can be applied to matrix product states, state vectors, quantum circuits, or any other framework in which a quantum state can be represented.

Our results demonstrate the feasibility of using the combination of the shift-invert method and imaginary time evolution to construct weakly entangled excited eigenstates of large quantum systems on classical hardware. We have demonstrated the method using classical tensor network techniques; however, as with all classical methods, the computational cost will grow rapidly as the system becomes more entangled. Generic midspectrum eigenstates are expected to

exhibit volume-law entanglement, rendering tensor network approaches inefficient in the worst case. A quantum circuit can capture such states without incurring exponential classical memory or bond-dimension overhead, so that the principal costs shift to state preparation and measurement. In this sense, our approach differs from algorithms specific to tensor networks, such as DMRG-X or shift and invert matrix product states (SIMPS). Our method directly targets eigenstates near a chosen energy via a spectral transformation and is not tied to any specific representation of the state, making it naturally compatible with quantum circuit-based implementations in regimes where classical descriptions become inefficient.

In this setting, one would variationally optimize the parameters of a parametrized quantum circuit, but otherwise the key steps of the algorithm would be precisely the same. We leave this as a goal for future work, as our focus in this Letter has been to develop and benchmark the algorithm itself, allowing us to illustrate the fundamental mechanism underlying the technique without encountering noise or precision issues. Nonetheless, such issues are important for practical implementations, and further details regarding scaling and resource estimation may be found in Ref. [28]. One particularly promising near-term hybrid approach would be to initially run the algorithm classically and then map the resulting MPS onto a quantum circuit using established techniques [38–41], which could then be used as a “warm start” initial configuration for the algorithm on quantum hardware. Alternatively, for translationally invariant systems, one could implement the algorithm using the infinite MPS (iMPS) [52] or infinite projected entangled pair states (PEPS) [53] frameworks, where only a single unit cell is considered. Similar strategies exploiting translation invariance can be employed for quantum circuits. This would significantly lower the computational cost, as only a single unit cell would need to be optimized, potentially allowing construction of highly excited, highly entangled states in one and two dimensions.

Our focus here has been on highly excited states; however, the method is very general and can also be used to target lower-lying excited states. For example, one could imagine finding the ground-state energy of a gapped system using, e.g., DMRG [18,19], VQE [36,37], or sample-based quantum diagonalization [54–56], and then constructing the first excited state using our method, enabling an accurate determination of the energy gap. Likewise, given an approximate location of the band gap in a material, our technique could be used to construct the eigenstates immediately above and below the

band gap, along with their eigenvalues, enabling the band gap itself to be precisely measured. In fact, we need not restrict ourselves to excited states: We can also use this method to directly target ground states. The main reason to do so is that if δ can be chosen to be sufficiently close to the ground-state energy, the resulting $\mathcal{H}_{\text{eff}}$ will have a larger energy gap than the initial Hamiltonian [28]. This may make it easier for variational methods to converge in systems where the initial Hamiltonian is gapless or otherwise has a challenging energy landscape. Classically efficient lower bounds for ground-state energies may be used to provide good choices for δ in this case [57–59].

Similar variational methods have been widely used in many proposed near-term quantum algorithms [60,61], including in the case of imaginary time evolution [62,63]. However, they have not to our knowledge been used in conjunction with the shift-invert technique. To that end, while this algorithm will likely be more widely used on classical hardware in the near term, we anticipate that its true long-term value will be in constructing highly entangled states directly on quantum hardware.

Acknowledgments. We acknowledge helpful discussions with L. Henaut, C. Willby, N. Mariella, T. Murphy, and O. T. Brown. S.J.T. thanks J. Eisert for pointing out the existence of efficiently computable lower bounds for ground-state energies. This work was supported by the Hartree National Centre for Digital Innovation, a UK Government-funded collaboration between STFC and IBM, as well as the Engineering and Physical Sciences Research Council (Grant No. EP/Z533518/1). All numerical results were obtained on the University of Edinburgh School of Physics & Astronomy Compute Cluster.

IBM, the IBM logo, and [64] are trademarks of International Business Machines Corp., registered in many jurisdictions worldwide. Other product and service names might be trademarks of IBM or other companies. The current list of IBM trademarks is available in Ref. [69].

S.J.T. and D.A.M. contributed equally to this work. S.J.T. developed the core algorithm. S.J.T. and D.A.M. jointly wrote the MPS code and performed High Performance Computing (HPC) simulations. L.W.A. and D.A.M. made key contributions to ensuring feasibility on quantum hardware. M.E.S. assisted with running high performance computing (HPC) simulations. All authors contributed to writing the final manuscript.

Data availability. The code and data that support the findings of this article are openly available [70].

- [1] S. Mishra, G. Catarina, F. Wu, R. Ortiz, D. Jacob, K. Eimre, J. Ma, C. A. Pignedoli, X. Feng, P. Ruffieux, *et al.*, Observation of fractional edge excitations in nanographene spin chains, *Nature (London)* **598**, 287 (2021).
- [2] A. K. Bera, S. Yusuf, S. K. Saha, M. Kumar, D. Voneshen, Y. Skourski, and S. A. Zvyagin, Emergent many-body composite excitations of interacting spin-1/2 trimers, *Nat. Commun.* **13**, 6888 (2022).
- [3] S. Bayliss, D. Laorenza, P. Mintun, B. Kovos, D. E. Freedman, and D. Awschalom, Optically addressable molecular spins for quantum information processing, *Science* **370**, 1309 (2020).
- [4] K. Katcko, E. Urbain, F. Ngassam, L. Kandpal, B. Chowrira, F. Schleicher, U. Halisdemir, D. Wang, T. Scherer, D. Mertz, *et al.*, Encoding information on the excited state of a molecular spin chain, *Adv. Funct. Mater.* **31**, 2009467 (2021).
- [5] D. Cocchiarella and M. C. Bañuls, Low-temperature Gibbs states with tensor networks, *Phys. Rev. Res.* **7**, 033184 (2025).
- [6] J. Jankowska and A. L. Sobolewski, Modern theoretical approaches to modeling the excited-state intramolecular proton transfer: An overview, *Molecules* **26**, 5140 (2021).

- [7] D. Pfau, S. Axelrod, H. Sutterud, I. von Glehn, and J. S. Spencer, Accurate computation of quantum excited states with neural networks, *Science* **385**, eadn0137 (2024).
- [8] V. Balzani, P. Ceroni, and A. Juris, *Photochemistry and Physics: Concepts, Research, Applications* (John Wiley & Sons, US, 2014).
- [9] P. O. Dral and M. Barbatti, Molecular excited states through a machine learning lens, *Nat. Rev. Chem.* **5**, 388 (2021).
- [10] J. Westermayr and P. Marquetand, Machine learning for electronically excited states of molecules, *Chem. Rev.* **121**, 9873 (2021).
- [11] A. Zewail, *The Chemical Bond: Structure and Dynamics* (Elsevier, US, 1992).
- [12] C. E. Crespo-Hernández, B. Cohen, P. M. Hare, and B. Kohler, Ultrafast excited-state dynamics in nucleic acids, *Chem. Rev.* **104**, 1977 (2004).
- [13] G. E. Crisenza, D. Mazzarella, and P. Melchiorre, Synthetic methods driven by the photoactivity of electron donor–acceptor complexes, *J. Am. Chem. Soc.* **142**, 5461 (2020).
- [14] A. Stoneham, Exploiting the excited state, *Physica B* **340–342**, 48 (2003).
- [15] V. W.-W. Yam, Using synthesis to steer excited states and their properties and functions, *Nat. Synth.* **2**, 94 (2023).
- [16] B. Hammond, W. Lester, and P. Reynolds, *Monte Carlo Methods in Ab Initio Quantum Chemistry*, Lecture and Course Notes In Chemistry Series (World Scientific, Singapore, 1994).
- [17] W. M. C. Foulkes, L. Mitas, R. J. Needs, and G. Rajagopal, Quantum Monte Carlo simulations of solids, *Rev. Mod. Phys.* **73**, 33 (2001).
- [18] U. Schollwöck, The density-matrix renormalization group in the age of matrix product states, *Ann. Phys.* **326**, 96 (2011).
- [19] J. C. Bridgeman and C. T. Chubb, Hand-waving and interpretive dance: An introductory course on tensor networks, *J. Phys. A: Math. Theor.* **50**, 223001 (2017).
- [20] F. Alet and N. Laflorencie, Many-body localization: An introduction and selected topics, *C. R. Phys.* **19**, 498 (2018).
- [21] D. A. Abanin, E. Altman, I. Bloch, and M. Serbyn, *Colloquium: Many-body localization, thermalization, and entanglement*, *Rev. Mod. Phys.* **91**, 021001 (2019).
- [22] P. Sierant, M. Lewenstein, A. Scardicchio, L. Vidmar, and J. Zakrzewski, Many-body localization in the age of classical computing*, *Rep. Prog. Phys.* **88**, 026502 (2025).
- [23] V. Khemani, F. Pollmann, and S. L. Sondhi, Obtaining highly excited eigenstates of many-body localized hamiltonians by the density matrix renormalization group approach, *Phys. Rev. Lett.* **116**, 247204 (2016).
- [24] F. Pollmann, V. Khemani, J. I. Cirac, and S. L. Sondhi, Efficient variational diagonalization of fully many-body localized Hamiltonians, *Phys. Rev. B* **94**, 041116(R) (2016).
- [25] X. Yu, D. Pekker, and B. K. Clark, Finding matrix product state representations of highly excited eigenstates of many-body localized hamiltonians, *Phys. Rev. Lett.* **118**, 017201 (2017).
- [26] S.-Y. Zhang, D. Yuan, T. Iadecola, S. Xu, and D.-L. Deng, Extracting quantum many-body scarred eigenstates with matrix product states, *Phys. Rev. Lett.* **131**, 020402 (2023).
- [27] L. Cadi Tazi and A. J. W. Thom, Folded Spectrum VQE: A quantum computing method for the calculation of molecular excited states, *J. Chem. Theory Comput.* **20**, 2491 (2024).
- [28] See Supplemental Material at <http://link.aps.org/supplemental/10.1103/v4rc-gmcx> for details of the quantum computing implementation, resource estimation, details of the MPS implementation, and additional classical simulations, which contains Refs. [65–68].
- [29] F. Pietracaprina, N. Macé, D. J. Luitz, and F. Alet, Shift-invert diagonalization of large many-body localizing spin chains, *SciPost Phys.* **5**, 045 (2018).
- [30] A. Altland and B. D. Simons, *Condensed Matter Field Theory*, 2nd ed. (Cambridge University Press, UK, 2010).
- [31] L. Cincio, Y. Subaşı, A. T. Sornborger, and P. J. Coles, Learning the quantum algorithm for state overlap, *New J. Phys.* **20**, 113022 (2018).
- [32] C. Cao, H. Yano, and Y. O. Nakagawa, Accelerated variational quantum eigensolver with joint Bell measurement, *Phys. Rev. Res.* **6**, 013205 (2024).
- [33] R. Puig, M. Drudis, S. Thanasilp, and Z. Holmes, Variational quantum simulation: A case study for understanding warm starts, *PRX Quantum* **6**, 010317 (2025).
- [34] G. Brassard, P. Høyer, M. Mosca, and A. Tapp, *Quantum Amplitude Amplification and Estimation*, Quantum Computation and Information, Contemporary Mathematics (2000), Vol. 305, pp. 53–74.
- [35] D. W. Berry, A. M. Childs, R. Cleve, R. Kothari, and R. D. Somma, Exponential improvement in precision for simulating sparse Hamiltonians, in *Proceedings of the Forty-Sixth Annual ACM Symposium on Theory of Computing* (Association for Computing Machinery, New York, 2014), pp. 283–292.
- [36] 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).
- [37] J. Tilly, H. Chen, S. Cao, D. Picozzi, K. Setia, Y. Li, E. Grant, L. Wossnig, I. Rungger, G. H. Booth, and J. Tennyson, The variational quantum eigensolver: A review of methods and best practices, *Phys. Rep.* **986**, 1 (2022).
- [38] S.-J. Ran, Encoding of matrix product states into quantum circuits of one- and two-qubit gates, *Phys. Rev. A* **101**, 032310 (2020).
- [39] D. Malz, G. Styliaris, Z.-Y. Wei, and J. I. Cirac, Preparation of matrix product states with log-depth quantum circuits, *Phys. Rev. Lett.* **132**, 040404 (2024).
- [40] D. A. Millar, G. W. Pennington, N. T. M. Siow, and S. J. Thomson, Qiskit-community/mps-to-circuit, 2025, <https://github.com/qiskit-community/mps-to-circuit>.
- [41] N. Robertson, A. Akhriev, J. Vala, and S. Zhuk, Approximate quantum compiling for quantum simulation: A tensor network based approach, *ACM Trans. Quantum Comput.* **6**, 1 (2025).
- [42] M. Fishman, S. R. White, and E. M. Stoudenmire, The ITensor Software Library for Tensor Network Calculations, *SciPost Phys. Codebases* **4** (2022).
- [43] K. S. Rai, J. I. Cirac, and Á. M. Alhambra, Matrix product state approximations to quantum states of low energy variance, *Quantum* **8**, 1401 (2024).
- [44] S.-H. Lin, R. Dilip, A. G. Green, A. Smith, and F. Pollmann, Real- and imaginary-time evolution with compressed quantum circuits, *PRX Quantum* **2**, 010342 (2021).
- [45] D. J. Luitz, N. Laflorencie, and F. Alet, Many-body localization edge in the random-field Heisenberg chain, *Phys. Rev. B* **91**, 081103(R) (2015).

- [46] S. J. Thomson, Localization and subdiffusive transport in quantum spin chains with dilute disorder, *Phys. Rev. B* **107**, 014207 (2023).
- [47] E. V. H. Doggen, F. Schindler, K. S. Tikhonov, A. D. Mirlin, T. Neupert, D. G. Polyakov, and I. V. Gornyi, Many-body localization and delocalization in large quantum chains, *Phys. Rev. B* **98**, 174202 (2018).
- [48] J. Šuntajs, J. Bonča, T. Prosen, and L. Vidmar, Quantum chaos challenges many-body localization, *Phys. Rev. E* **102**, 062144 (2020).
- [49] D. Sels and A. Polkovnikov, Dynamical obstruction to localization in a disordered spin chain, *Phys. Rev. E* **104**, 054105 (2021).
- [50] D. Sels and A. Polkovnikov, Thermalization of dilute impurities in one-dimensional spin chains, *Phys. Rev. X* **13**, 011041 (2023).
- [51] T. Devakul, V. Khemani, F. Pollmann, D. A. Huse, and S. L. Sondhi, Obtaining highly excited eigenstates of the localized XX chain via DMRG-X, *Philos. Trans. R. Soc. A* **375**, 20160431 (2017).
- [52] G. Vidal, Classical simulation of infinite-size quantum lattice systems in one spatial dimension, *Phys. Rev. Lett.* **98**, 070201 (2007).
- [53] J. Jordan, R. Orús, G. Vidal, F. Verstraete, and J. I. Cirac, Classical simulation of infinite-size quantum lattice systems in two spatial dimensions, *Phys. Rev. Lett.* **101**, 250602 (2008).
- [54] K. Kanno, M. Kohda, R. Imai, S. Koh, K. Mitarai, W. Mizukami, and Y. O. Nakagawa, Quantum-selected configuration interaction: Classical diagonalization of Hamiltonians in subspaces selected by quantum computers, [arXiv:2302.11320](https://arxiv.org/abs/2302.11320).
- [55] J. Robledo-Moreno, M. Motta, H. Haas, A. Javadi-Abhari, P. Jurcevic, W. Kirby, S. Martiel, K. Sharma, S. Sharma, T. Shirakawa, *et al.*, Chemistry beyond the scale of exact diagonalization on a quantum-centric supercomputer, *Sci. Adv.* **11**, eadu9991 (2025).
- [56] S. Barison, J. R. Moreno, and M. Motta, Quantum-centric computation of molecular excited states with extended sample-based quantum diagonalization, *Quantum Sci. Technol.* **10**, 025034 (2025).
- [57] P. W. Anderson, Limits on the energy of the antiferromagnetic ground state, *Phys. Rev.* **83**, 1260 (1951).
- [58] T. Baumgratz and M. B. Plenio, Lower bounds for ground states of condensed matter systems, *New J. Phys.* **14**, 023027 (2012).
- [59] J. Eisert, Lower bounds to variational problems with guarantees, *Phys. Rev. A* **113**, 022214 (2026).
- [60] X. Yuan, S. Endo, Q. Zhao, Y. Li, and S. C. Benjamin, Theory of variational quantum simulation, *Quantum* **3**, 191 (2019).
- [61] M. Cerezo, A. Arrasmith, R. Babbush, S. C. Benjamin, S. Endo, K. Fujii, J. R. McClean, K. Mitarai, X. Yuan, L. Cincio, and P. J. Coles, Variational quantum algorithms, *Nat. Rev. Phys.* **3**, 625 (2021).
- [62] 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).
- [63] M. Motta, C. Sun, A. 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).
- [64] www.ibm.com.
- [65] H. Chernoff, A measure of asymptotic efficiency for tests of a hypothesis based on the sum of observations, *Ann. Math. Stat.* **23**, 493 (1952).
- [66] W. Hoeffding, Probability inequalities for sums of bounded random variables, *J. Am. Stat. Assoc.* **58**, 13 (1963).
- [67] C. Hubig, I. P. McCulloch, U. Schollwöck, and F. A. Wolf, Strictly single-site DMRG algorithm with subspace expansion, *Phys. Rev. B* **91**, 155115 (2015).
- [68] D. Wierichs, J. Izaac, C. Wang, and C. Y.-Y. Lin, General parameter-shift rules for quantum gradients, *Quantum* **6**, 677 (2022).
- [69] www.ibm.com/legal/copytrade.
- [70] D. A. Millar, L. W. Anderson, E. Altamura, O. Wallis, E. Sahin, J. Crain, and S. J. Thomson, Data and code underpinning “Imaginary time spectral transforms for excited state preparation”, Zenodo, 2025, <https://zenodo.org/records/16635254>.