

Large-Scale Implementation of Quantum Subspace Expansion with Classical Shadows

Laurin E. Fischer^{1,2,*}, Daniel Bultrini³, Timothée Dao^{1,4}, Ivano Tavernelli¹, and Francesco Tacchino^{1,†,‡}
¹IBM Research, 8803 Rüschlikon, Switzerland

²Theory and Simulation of Materials, École Polytechnique Fédérale de Lausanne, 1015 Lausanne, Switzerland

³Theoretische Chemie, Physikalisch-Chemisches Institut, Heidelberg University, Heidelberg, Germany

⁴Institute for Theoretical Physics, ETH Zürich, 8093 Zürich, Switzerland

 (Received 4 December 2025; revised 22 May 2026; accepted 24 August 2026; published 15 October 2026)

Quantum subspace expansion (QSE) offers promising avenues to perform spectral calculations on quantum processors but comes with a large measurement overhead. Informationally complete measurements, such as classical shadows, were recently proposed to overcome this bottleneck. Here, we report the first large-scale implementation of QSE with informationally complete measurements. In particular, we probe the quantum phase transition of a spin model with three-body interactions, for which we observe accurate ground state energy recovery and mitigation of local order parameters across system sizes of up to 80 qubits. We achieve this by reformulating QSE as a constrained optimization problem, obtaining rigorous statistical error estimates and avoiding numerical ill conditioning. With over 3×10^4 measurement basis randomizations per circuit and the evaluation of $\mathcal{O}(10^{14})$ Pauli traces, this represents one of the most significant realizations of classical shadows to date.

DOI: [10.1103/2w21-qs38](https://doi.org/10.1103/2w21-qs38)

Introduction—Spectral calculations are a key target for many-body quantum simulation frameworks. Until recently, heuristic approaches based on the *variational quantum eigensolver* dominated the near-term quantum computing literature on the topic, thanks to their relatively low demands on quantum processor performance [1]. However, practical issues connected to trainability [2] and large measurement overheads [3] have so far hindered their scalability beyond small proof-of-principle demonstrations. In response to this, alternative methods for calculating ground and excited states are being explored. A promising class of post-variational-quantum-eigensolver quantum algorithms for spectral calculations is *quantum subspace expansion* (QSE) [4]. Assume that, as usual, the goal is to find the energy E_{gs} of a ground state $|\psi_{\text{gs}}\rangle$ of a target Hamiltonian H , and suppose that a certain “root state,” ρ_0 , close to $|\psi_{\text{gs}}\rangle$ can be prepared on a quantum computer. QSE relies on the fact that, while ρ_0 alone may not yield the desired accuracy, e.g., due to noise or unitary circuit approximations, it might still represent a useful resource. In fact, there exist several ways in which ρ_0 can be employed to inform a subsequent classical calculation.

These include, for instance, matrix diagonalization based on projection into a space spanned by individual samples [5,6], the extraction of an effective representation of ρ_0 based on classical shadows [7], and subsequent processing of the quantum computer outputs via tensor networks [8,9] or neural quantum states [10].

In QSE, a space of states is spanned on top of ρ_0 with L Hermitian expansion operators σ_i as [11] $\rho_{\text{QSE}}(\vec{c}) = W^\dagger \rho_0 W / \text{Tr}[W^\dagger \rho_0 W]$, with $W = \sum_{i=1}^L c_i \sigma_i$, parametrized by a coefficient vector $\vec{c} = (c_1, \dots, c_L)$. This can be regarded as accessing additional states beyond the scope of the quantum processor [12,13], and may also serve as noise-agnostic error mitigation [11,14–16]. The expectation value of an observable O over a state in the subspace then becomes

$$\text{Tr}[O \rho_{\text{QSE}}(\vec{c})] = \frac{\sum_{i,j=1}^L c_i^* c_j \mathcal{O}_{ij}}{\sum_{i,j=1}^L c_i^* c_j \mathcal{S}_{ij}}, \quad (1)$$

where $\mathcal{O}_{ij} = \text{Tr}[\sigma_i^\dagger \rho_0 \sigma_j O]$ and $\mathcal{S}_{ij} = \text{Tr}[\sigma_i^\dagger \rho_0 \sigma_j]$. For ground state search, the goal is to find the minimal expectation value $\min_{\vec{c}} \text{Tr}[H \rho_{\text{QSE}}(\vec{c})]$, which corresponds to the smallest pseudoeigenvalue λ of the generalized eigenvalue problem $\mathcal{H}\vec{c} = \lambda \mathcal{S}\vec{c}$, where $\mathcal{H}$ is the subspace-projected Hamiltonian and $\mathcal{S}$ is called the overlap matrix. For reasonable values of L , this is tractable on a classical computer by inverting the overlap matrix and computing the smallest eigenvalue of $\mathcal{S}^{-1}\mathcal{H}$.

The main measurement cost of QSE stems from the number of measurements required to estimate all matrix entries $\mathcal{O}_{ij}$ and $\mathcal{S}_{ij}$, which depends on the dimension L and

*Contact author: aur@zurich.ibm.com

†Contact author: francesco.tacchino@unipv.it

‡Present address: Università di Pavia, Pavia, Italy.

Published by the American Physical Society under the terms of the [Creative Commons Attribution 4.0 International](https://creativecommons.org/licenses/by/4.0/) license. Further distribution of this work must maintain attribution to the author(s) and the published article's title, journal citation, and DOI.

the complexity of the expansion operators. For approaches that measure each observable individually (or in commuting groups) this can be a significant roadblock for scaling in practice [17,18]. To alleviate this, *informationally complete* (IC) measurements, such as *classical shadows* [7] (CSs) can be used to estimate all underlying observables from one set of measurement samples [19,20]. CS measurements have recently been realized in a series of large-scale quantum simulations [21,22] and the use of CSs to systematically build high-dimensional subspaces spanned by low-weight Paulis has been proposed in Ref. [23]. However, while the latter offered a promising avenue, its statistical analysis focused on worst-case upper bounds rather than reliable statistical error bars for the resulting estimators. Furthermore, no demonstration of such ideas on quantum hardware has been achieved so far. In this Letter, we close both gaps by formulating subspace expansion as a constrained optimization problem. In contrast to previous methods, this allows us to obtain accurate statistical uncertainties and provides an external tuning knob to control these errors, trading potential bias for variance. With this technique, we demonstrate noise-agnostic error mitigation via QSE in ground states of spin models for systems with up to 80 qubits using state preparations circuits introduced in Refs. [24,25].

Theory—An IC measurement is described by a set of n positive semidefinite Hermitian operators $\{M_k\}_{k \in \{1, \dots, n\}}$ with $\sum_{k=1}^n M_k = \mathbb{1}$ such that any observable O can be expressed as $O = \sum_{k=1}^n \omega_k M_k$ for some $\omega_k \in \mathbb{R}$ [26]. The expectation value $\langle O \rangle$ then becomes $\text{Tr}[\rho_0 O] = \sum_k \omega_k \text{Tr}[\rho_0 M_k] = \mathbb{E}_{k \sim p_k}[\omega_k]$, i.e., $\langle O \rangle$ can be expressed as the mean value of the random variable ω_k over the probability distribution $\{p_k = \text{Tr}[\rho_0 M_k]\}$. Given a sample of S measurement outcomes $\{k^{(1)}, \dots, k^{(S)}\}$, we can thus construct an unbiased estimator of $\langle O \rangle$ as

$$\hat{o} : \{k^{(1)}, \dots, k^{(S)}\} \mapsto \frac{1}{S} \sum_{s=1}^S \omega_{k^{(s)}}. \quad (2)$$

The statistical variance of this estimator is given by the standard error on the mean $\text{Var}[\hat{o}] = \text{Var}[\omega_k]/S$. The coefficients ω_k can be obtained from frame theory [27] where each positive operator-valued measure (POVM) operator M_k is associated with a dual operator D_k (also known as the “classical shadow” of a state) such that $\omega_k = \text{Tr}[OD_k]$.

Given S POVM outcome samples $\{k^{(1)}, \dots, k^{(S)}\}$ measured from the base state ρ_0 , we aim to simultaneously estimate all matrix entries of the projected Hamiltonian $\mathcal{H}_{ij}$ and the overlap matrix $\mathcal{S}_{ij}$. Following Eq. (2), a canonical estimator for each matrix element $\mathcal{H}_{ij}$ and $\mathcal{S}_{ij}$ is obtained by simply replacing ρ_0 in each expression with the corresponding dual operator D_k and then averaging over all shots. In principle we could estimate all entries of $\mathcal{H}$ and

$\mathcal{S}$ this way and then solve the generalized eigenvalue problem. However, in practice this direct inversion of the overlap matrix can be numerically unstable due to ill conditioning of the matrix under shot noise and hardware imperfections, particularly when the expansion operators are nonorthogonal [28]. This issue is often addressed through regularization of the overlap matrix, e.g., by discarding dimensions that correspond to small singular values of $\mathcal{S}$ [13]. Furthermore, since with IC measurements we do not estimate the matrix entries independently, additional covariances between them need to be accounted for when deriving statistical error bars.

Here, we present a different scheme to avoid ill conditioning that is tailored to handle such inherent covariances. Assume that we have an estimator $\hat{H}(\vec{c})$ of the energy $\text{Tr}[H\rho_{\text{QSE}}(\vec{c})]$ of the state at subspace vector $\vec{c}$ and that we can further estimate the statistical error of $\hat{H}(\vec{c})$ as $\hat{\epsilon}(\vec{c})$. In practice, we would like to compute the lowest possible subspace energy given a maximally tolerated error ϵ_{max} . This is achieved through the constrained optimization problem

$$\arg \min_{\vec{c} \in \mathbb{R}^L} \hat{H}(\vec{c}) \text{ subject to } \hat{\epsilon}(\vec{c}) \leq \epsilon_{\text{max}}. \quad (3)$$

This allows us to explore only those regions of the subspace where the statistical error is well controlled. In practice, the solution of Eq. (3) typically saturates the constraint $\hat{\epsilon}(\vec{c}) \leq \epsilon_{\text{max}}$, such that increasing ϵ_{max} enables a smooth exploration of lower-energy subspace states at the cost of increased statistical uncertainty. In contrast to regularization techniques, this provides a direct control of the bias-variance trade-off inherent in the estimator; see Fig. 1 and Supplemental Material [29].




FIG. 1. Subspace expansion as a constrained optimization problem. The z axis shows the estimated energy $\hat{H}(\vec{c})$ when varying two dimensions of the subspace coefficients $\vec{c}$, while color indicates the statistical error $\hat{\epsilon}(\vec{c})$ of the energy estimation. Naive minimization of the energy (red arrow) may yield an energy with high statistical error that can severely violate the variational principle (black circle indicates the true ground state energy). Constrained optimization with a maximum allowed error ϵ_{max} avoids statistically unstable regions (blue path). Data corresponds to the results presented in Fig. 2 for $N = 48$ and $g = -0.5$.

We now construct the estimators $\hat{H}(\vec{c})$ and $\hat{e}(\vec{c})$ based on IC samples. Let $x_k = \sum_{i,j=1}^L c_i^* c_j \text{Tr}[D_k \sigma_j H \sigma_i^\dagger]$ and $y_k = \sum_{i,j=1}^L c_i^* c_j \text{Tr}[D_k \sigma_j \sigma_i^\dagger]$ be the random variables of the numerator and denominator of Eq. (1) distributed according to p_k . Then the subspace expectation value $\text{Tr}[H \rho_{\text{QSE}}(\vec{c})] = \mu_x / \mu_y$ is given by the ratio of $\mu_x = \mathbb{E}_{k \sim p_k}[x_k]$ and $\mu_y = \mathbb{E}_{k \sim p_k}[y_k]$. An asymptotically unbiased and consistent estimator of this ratio is given by the ratio of the sample means [36],

$$\hat{H}(\vec{c}) : \{k^{(1)}, \dots, k^{(S)}\} \mapsto \frac{\bar{x}}{\bar{y}}, \quad (4)$$

with $\bar{x} = \sum_{i=1}^S x_{k^{(i)}} / S$ (similarly for $\bar{y}$). While no closed-form expression exists for the variance $e^2(\vec{c}) = \text{Var}[\hat{H}(\vec{c})]$ of this estimator, it can be well approximated [37] by a second-order Taylor approximation around (μ_x, μ_y) , which yields [36]

$$e^2(\vec{c}) \approx \frac{1}{S} \left(\frac{\text{Var}[x_k]}{\mu_y^2} + \frac{\mu_x^2 \text{Var}[y_k]}{\mu_y^4} - \frac{2\mu_x \text{Cov}[x_k, y_k]}{\mu_y^3} \right). \quad (5)$$

Note that for $\sigma_0 = \mathbb{1}$ and $\vec{c} = (1, 0, 0, \dots)$ this reduces to the standard error on the mean. We obtain the estimated error $\hat{e}(\vec{c})$ by inserting the sample means $\bar{x}$ for μ_x , the sample standard deviations $\sum_{i=1}^S |x_{k^{(i)}} - \bar{x}|^2 / (S-1)$ for $\text{Var}[x_k]$ (similarly for y_k), and the sample covariance $\sum_{i=1}^S (x_{k^{(i)}} - \bar{x})(y_{k^{(i)}} - \bar{y}) / (S-1)$ for $\text{Cov}[x_k, y_k]$ into Eq. (5). The covariance term captures the contributions to the statistical error that arise from reusing the same measurement samples to estimate all matrix entries.

The above statistical error propagation holds when the POVM samples are independent and identically distributed (i.i.d.). In practice, for POVMs from randomized measurements (such as CS), it is often practically convenient to take multiple measurement shots for each randomized readout basis, which violates the i.i.d. assumption. In this case, we can redefine the random variables x_k and y_k as their averaged value over all repeated shots taken in the same measurement basis. In this way, the obtained values for x_k and y_k remain i.i.d., and Eqs. (4), (5) remain valid for solving the constrained optimization problem from Eq. (3).

Besides substantial efficiency gains from evaluating the matrix entries of the subspace expansion in parallel rather than sequentially [20], our approach based on IC measurements offers two additional advantages over traditional methods. Firstly, the expansion operators $\{\sigma_1, \dots, \sigma_L\}$ can be chosen *a posteriori* after measurements have been taken, as also pointed out by Ref. [23]. Thus, the same measurements can be reused to iteratively optimize the chosen expansion operators to minimize the energy while the constrained optimization guarantees that the variational principle is respected up to the chosen statistical precision ϵ_{max} . Secondly, the framework offers the flexibility of

estimating additional observables O for the obtained energy-optimized state without additional measurements by constructing the estimator $\hat{O}(\vec{c}_{\text{opt}})$ equivalently to Eq. (4) (replacing all occurrences of H with O), where $\vec{c}_{\text{opt}}$ is the solution to Eq. (3).

Hardware demonstration—We demonstrate our method by addressing the task of recovering the ground state energy of a Hamiltonian with known approximate ground state preparation circuits. We consider a one-dimensional spin chain originally studied in Ref. [38] described by the Hamiltonian (with periodic boundary conditions),

$$H = \sum_{i=1}^N (-g_{zz} Z_i Z_{i+1} - g_x X_i + g_{zxz} Z_i X_{i+1} Z_{i+2}), \quad (6)$$

where the parameter trajectory (for $g \in [-1, +1]$) $(g_{zz}(g), g_x(g), g_{zxz}(g)) = (2(1-g^2), (1+g)^2, (g-1)^2)$ traverses a quantum phase transition with a critical point at $g = 0$ between a symmetry-protected topological (SPT) phase where the ZZ terms dominate, and a trivial phase where the X terms dominate. Along this trajectory, the ground state of the model corresponds to a translation-invariant normal matrix product state (MPS) [39] with an energy density $E_{\text{gs}}/N = -2(g^2 + 1)$. We implement the latter as a quantum circuit leveraging a recently introduced algorithm based on the renormalization-group transformation that approximately prepares the MPS with an error δ in depth $\mathcal{O}(\log(N/\delta))$ [24]. The approximation is determined by a “blocking number” q , which navigates a trade-off between circuit depth and the fidelity of the prepared state. Here, we opt for a blocking number of $q = 4$, which has been shown to be a reasonable choice for current noisy quantum hardware [25].

We implement the MPS ansatz circuits with system sizes increasing from $N = 16$ with 108 entangling CZ gates to $N = 80$ with 540 CZ gates on the IBM quantum superconducting qubit device *ibm_fez* of the Heron R2 generation; see Supplemental Material [29] for details. We measure an IC POVM that consists of uniformly randomized single-qubit X, Y, and Z measurements (standard local classical shadows with canonical duals [27]; see also Supplemental Material [29]). For each studied system size N and parameter value g , we perform measurements in 32768 bases with eight shots per basis. This POVM can efficiently estimate observables of low Pauli weight w with statistical errors increasing exponentially with w [7]. Hence, in principle, we could now choose any set of sufficiently local observables as expansion operators σ_i of the subspace. We first select $\sigma_1 = \mathbb{1}$ to ensure that the original state ρ_0 is part of the subspace. In this way, the subspace-optimized energy can never be worse than the measured one.

For ground state search, the powers of the Hamiltonian H^p are popular expansion operators as the overlap of the states $H^p \rho_0 H^{p\dagger}$ with the true ground state increases

exponentially with p , provided ρ_0 has nonzero overlap with $|\psi_{\text{gs}}\rangle$. This corresponds to the well-known Krylov subspace diagonalization method [16,40–42]. We thus choose $\sigma_2 = H$ to boost the ground state population. However, in our model $\sigma_2 = H$ already leads to a term $\mathcal{H}_{2,2} = \text{Tr}[\rho_0 H^3]$, which is an observable with 2.14×10^6 Pauli terms at $N = 80$. The trace of each of these Pauli terms with each measured dual operator (which are 262144 in our case) needs to be evaluated. The inclusion of $\sigma_3 = H^2$ would result in a term $\mathcal{H}_{3,3} = \text{Tr}[\rho_0 H^5]$ with 35×10^9 Pauli terms and strings with weight up to $w = 15$. These both exceed our available classical processing resources and are beyond the scope of estimation with single-qubit classical shadows. Because of this scalability bottleneck (see Supplemental Material [29]), we limit ourselves to one Krylov dimension and label such expansion sets as “Krylov” = $\{1, H\}$.

Instead of increasing the Krylov dimension, we further lower the energy by systematically adding low-weight Pauli expansion operators with a strategy proposed similarly in Ref. [23]. For Pauli weights $w \in \{1, 2, 3, 4, 5\}$, we sample 150 Pauli strings of weight w uniformly at random

to form a pool $\mathcal{P}_w$. For each $P_i \in \mathcal{P}_w$, we perform a two-dimensional subspace expansion with operators $\{1, P_i\}$. We then rank the Pauli strings within each pool according to the amount by which they reduce the energy relative to the unmitigated estimate when optimizing for the lowest upper error given by $\arg \min_{\vec{c}} \hat{H}(\vec{c}) + \hat{\epsilon}(\vec{c})$. Finally, we include up to five of the best-performing Paulis for each considered weight into the set of expansion operators; see Supplemental Material [29]. Together with $\sigma_1 = 1$ and $\sigma_2 = H$, this results in a subspace that we refer to as “Krylov+.” The total numbers of single-qubit Pauli traces of the form $\text{Tr}, P_i, P_j \in \{I, X, Y, Z\}$ between dual operators of measurement outcomes and the observables of the subspace range from 4.3×10^{11} for $N = 16$ to 5.6×10^{13} for $N = 80$; see Supplemental Material [29]. For both considered subspaces, we solve Eq. (3) with the *constrained optimization by linear approximation* algorithm [43]. The allowed maximum errors ϵ_{max} were chosen as a fixed multiple of the unmitigated signal to ensure a signal-to-noise ratio ranging from 5% for $N = 16$ to 15% for $N = 80$; see Supplemental Material [29].




FIG. 2. (a) Error mitigation with subspace expansion based on IC measurements of ground state preparation circuits for a spin model. The qubit number increases from left to right from $N = 16$ to $N = 80$, while the three rows show the energy density and the order parameters of the phase transition traversed as a function of the Hamiltonian parameter g . Error bars represent the estimated statistical error $\hat{\epsilon}(\vec{c})$ and the insets show the chosen physical qubits of the 156-qubit device *ibm_fez*. For $N = 16$, gray shaded curves show the theoretical values for the lowest 25 excited states obtained with exact diagonalization, with a color gradient from black for the ground state to lighter shades for progressively higher-excited states. (b) Lowest pseudoeigenenergies obtained through regularization of the overlap matrix by discarding the lowest few singular values (SVs). Note that direct inversion of the problem without regularization leads to strong violations of the variational principle (gray region).

The results for the error-mitigated energy densities of the MPS ground state circuits for $g \in \{-0.9, -0.5, -0.1, 0.1, 0.5, 0.9\}$ are shown in the top row of Fig. 2(a). Across all studied system sizes, the Krylov subspace already significantly improves on the unmitigated energies, while the Krylov + approach manages to accurately recover the true ground state expectation values while respecting the variational principle within statistical uncertainties. As the system size increases, the statistical errors of the unmitigated results become progressively worse, which leads to larger error bars also for the mitigated values. For the smallest system size of $N = 16$ we can perform an exact diagonalization of the Hamiltonian to get the full low-energy (density) spectrum. This reveals that several states lie between the unmitigated and Krylov values and the true ground state. However, the Krylov + values lie well below the first excited state, indicating that the optimized subspace state indeed manages to produce significant overlap with the true ground state. We also compare to a direct solution of the generalized eigenvalue problem via inversion of the (regularized) overlap matrix. We find that this approach fails to safeguard against violations of the variational principle and does not recover the ground state energy accurately; see Fig. 2(b) and Supplemental Material [29]. Moreover, we validate the approximations that underlie the estimation of $\hat{e}(\vec{c})$ through bootstrap resampling and verify that they are accurate across most parameter regimes, with deviations arising primarily when denominator fluctuations are large; see Supplemental Material [29].

The phase transition at $g = 0$ is characterized by order parameters that serve as interesting additional observables for which we can reuse the IC measurement samples. While SPT phases in general are characterized by nonlocal order parameters, they can be expressed as local operators for the considered MPS ground states [39]. $S^X = \sum_{j=1}^N X_j/N$ signals the trivial phase. It is zero for states with $g < 0$ and goes to 1 with $g \rightarrow 1$ as $4g/(1+g)^2$. Similarly, the SPT phase is characterized by $S^{ZY} = \sum_{j=1}^N Z_j Y_{j+1} X_{j+2} Y_{j+3} Z_{j+4}/N$, which remains zero for $g > 0$ and goes to 1 with $g \rightarrow -1$ as $-4g/(1-g)^2$. The obtained values for S^X for the energy-optimized subspace state indeed improve over the unmitigated values for high g values; see second row of Fig. 2. However, the weight-five observable S^{ZY} shown in the final row of Fig. 2 is more difficult to mitigate. While the results recover some of the theoretical behavior for $N = 16$, the statistical uncertainties in the Krylov + values dominate at larger scales. From this, we conclude that, while the energy remains well-mitigated throughout all system sizes, the recovered subspace state does not provide sufficient overlap with the ground state to accurately reflect complex observables like S^{ZY} .

Discussion—In this Letter, we presented the first large-scale implementation of quantum subspace expansion

powered by classical shadows. Leveraging the high sampling rates of state-of-the-art superconducting quantum processors, we collected measurements in 32 768 randomized bases per input state and subsequently classically processed over 100×10^{14} individual Pauli traces. This not only represents the realization of classical shadow measurements at an unprecedented scale, but also offers a blueprint for future work flows combining quantum and high-performance classical computing.

Our results constitute a viable path to overcoming the measurement overhead of QSE, and include—as a key technical contribution—a complete characterization of its statistical error bars in the presence of covariances arising from the use of IC data. Furthermore, our results underscore the potential of QSE as a fully noise model-agnostic error mitigation technique [16], which, in contrast to other methods [21,44,45], does not rely on an accurate device noise characterization.

Our approach could be extended by incorporating symmetry verification [14], powers of the root state in the spirit of virtual distillation [37] and multireference methods [11]. Additionally, estimators based on medians-of-means [7] or dual frame optimization [46–48] could further reduce the sampling costs.

Overall, our framework provides a resource-aware approach that exposes and balances quantum coherence, sampling, and postprocessing costs. Importantly, the root state ρ_0 should be only a few, efficiently measurable expansion operators away from the target one. The number and nature of these operators are jointly determined by the available classical computational resources and the chosen family of IC measurements. For example, the single-qubit shadows used here efficiently support low-weight operators, whereas alternative families (e.g., shallow [49,50], fermionic [51,52], or global Clifford shadows [7]) extend to broader expansion sets. The constrained optimization and statistical error analyses carry over unchanged to any such setting, as illustrated for a molecular subspace expansion with fermionic shadows in Supplemental Material [29]. Interestingly, this probes a regime that is complementary to that of single-shot-based subspace expansion [5,6], whose success hinges on the fact that the ground state is only supported on polynomially many states of the sampled basis.

Acknowledgments—We thank Moritz Scheer, Alberto Baiardi, and Elisa Bäumer for support on constructing MPS ground state circuits. We thank Guillermo García-Pérez and Oriol Vendrell for helpful discussions in the early stages of the project. L. E. F. and D. B. acknowledge funding from the European Union’s Horizon 2020 research and innovation program under the Marie Skłodowska-Curie Grant Agreement No. 955479 (MOQS—Molecular Quantum Simulations). This research was supported by NCCR MARVEL, a National Center of Competence in

Research, funded by the Swiss National Science Foundation (Grant No. 205602) and by RESQUE, funded by the Swiss National Science Foundation (Grant No. 225229).

Data availability—The raw data that support the findings of this article are openly available at [53].

- [1] 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).
- [2] M. Larocca, S. Thanasilp, S. Wang, K. Sharma, J. Biamonte, P. J. Coles, L. Cincio, J. R. McClean, Z. Holmes, and M. Cerezo, Barren plateaus in variational quantum computing, *Nat. Rev. Phys.* **7**, 174 (2025).
- [3] J. F. Gonthier, M. D. Radin, C. Buda, E. J. Daskocil, C. M. Abuan, and J. Romero, Measurements as a roadblock to near-term practical quantum advantage in chemistry: Resource analysis, *Phys. Rev. Res.* **4**, 033154 (2022).
- [4] M. Motta, W. Kirby, I. Liepuoniute, K. J. Sung, J. Cohn, A. Mezzacapo, K. Klymko, N. Nguyen, N. Yoshioka, and J. E. Rice, Subspace methods for electronic structure simulations on quantum computers, *Electron. Struct.* **6**, 013001 (2024).
- [5] 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, *Phys. Rev. Res.* **8**, 023268 (2026).
- [6] 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).
- [7] H.-Y. Huang, R. Kueng, and J. Preskill, Predicting many properties of a quantum system from very few measurements, *Nat. Phys.* **16**, 1050 (2020).
- [8] G. García-Pérez, E.-M. Borrelli, M. Leahy, J. Malmi, S. Maniscalco, M. A. C. Rossi, B. Sokolov, and D. Cavalcanti, Virtual linear map algorithm for classical boost in near-term quantum computing, [arXiv:2207.01360](https://arxiv.org/abs/2207.01360).
- [9] S. Filippov, M. Leahy, M. A. Rossi, and G. García-Pérez, Scalable tensor-network error mitigation for near-term quantum computing, [arXiv:2307.11740](https://arxiv.org/abs/2307.11740).
- [10] S.-X. Zhang, Z.-Q. Wan, C.-K. Lee, C.-Y. Hsieh, S. Zhang, and H. Yao, Variational quantum-neural hybrid eigensolver, *Phys. Rev. Lett.* **128**, 120502 (2022).
- [11] N. Yoshioka, H. Hakoshima, Y. Matsuzaki, Y. Tokunaga, Y. Suzuki, and S. Endo, Generalized quantum subspace expansion, *Phys. Rev. Lett.* **129**, 020502 (2022).
- [12] T. Takeshita, N. C. Rubin, Z. Jiang, E. Lee, R. Babbush, and J. R. McClean, Increasing the representation accuracy of quantum simulations of chemistry without extra quantum resources, *Phys. Rev. X* **10**, 011004 (2020).
- [13] M. Urbanek, D. Camps, R. Van Beeumen, and W. A. De Jong, Chemistry on quantum computers with virtual quantum subspace expansion, *J. Chem. Theory Comput.* **16**, 5425 (2020).
- [14] X. Bonet-Monroig, R. Sagastizabal, M. Singh, and T. E. O'Brien, Low-cost error mitigation by symmetry verification, *Phys. Rev. A* **98**, 062339 (2018).
- [15] Z. Cai, Quantum error mitigation using symmetry expansion, *Quantum* **5**, 548 (2021).
- [16] P. Suchsland, F. Tacchino, M. H. Fischer, T. Neupert, P. K. Barkoutsos, and I. Tavernelli, Algorithmic error mitigation scheme for current quantum processors, *Quantum* **5**, 492 (2021).
- [17] P. J. Ollitrault, A. Kandala, C.-F. Chen, P. K. Barkoutsos, A. Mezzacapo, M. Pistoia, S. Sheldon, S. Wörner, J. M. Gambetta, and I. Tavernelli, Quantum equation of motion for computing molecular excitation energies on a noisy quantum processor, *Phys. Rev. Res.* **2**, 043140 (2020).
- [18] J. Selisko, M. Amsler, C. Wever, Y. Kawashima, G. Samsonidze, R. U. Haq, F. Tacchino, I. Tavernelli, and T. Eickl, Dynamical mean field theory for real materials on a quantum computer, [arXiv:2404.09527](https://arxiv.org/abs/2404.09527).
- [19] G. García-Pérez, M. A. Rossi, B. Sokolov, F. Tacchino, P. K. Barkoutsos, G. Mazzola, I. Tavernelli, and S. Maniscalco, Learning to measure: Adaptive informationally complete generalized measurements for quantum algorithms, *PRX Quantum* **2**, 040342 (2021).
- [20] S. Choi and A. F. Izmaylov, Measurement optimization techniques for excited electronic states in near-term quantum computing algorithms, *J. Chem. Theory Comput.* **19**, 3184 (2023).
- [21] L. E. Fischer, M. Leahy, A. Eddins, N. Keenan, D. Ferracin, M. A. Rossi, Y. Kim, A. He, F. Pietracaprina, B. Sokolov *et al.*, Dynamical simulations of many-body quantum chaos on a quantum computer, *Nat. Phys.* **1** (2026).
- [22] M. Votto, M. Ljubotina, C. Lancien, J. I. Cirac, P. Zoller, M. Serbyn, L. Piroli, and B. Vermersch, Learning mixed quantum states in large-scale experiments, *Phys. Rev. Lett.* **136**, 090801 (2026).
- [23] G. Boyd, B. Koczor, and Z. Cai, High-dimensional subspace expansion using classical shadows, *Phys. Rev. A* **111**, 022423 (2025).
- [24] 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).
- [25] M. Scheer, A. Baiardi, E. Bäumer Marty, Z.-Y. Wei, and D. Malz, Renormalization-group-based preparation of matrix product states on up to 80 qubits, [arXiv:2510.24681](https://arxiv.org/abs/2510.24681).
- [26] L. E. Fischer, D. Miller, F. Tacchino, P. K. Barkoutsos, D. J. Egger, and I. Tavernelli, Ancilla-free implementation of generalized measurements for qubits embedded in a qudit space, *Phys. Rev. Res.* **4**, 033027 (2022).
- [27] L. Innocenti, S. Lorenzo, I. Palmisano, F. Albarelli, A. Ferraro, M. Paternostro, and G. M. Palma, Shadow tomography on general measurement frames, *PRX Quantum* **4**, 040328 (2023).
- [28] E. N. Epperly, L. Lin, and Y. Nakatsukasa, A theory of quantum subspace diagonalization, *SIAM J. Matrix Anal. Appl.* **43**, 1263 (2022).
- [29] See Supplemental Material at <http://link.aps.org/supplemental/10.1103/2w21-qs38> for details on the hardware demonstration, classical postprocessing, bootstrap validation of statistical uncertainties, and an application to molecular systems with fermionic shadows, which includes Refs. [30–35].

- [30] A. Javadi-Abhari, M. Treinish, K. Krsulich, C. J. Wood, J. Lishman, J. Gacon, S. Martiel, P. D. Nation, L. S. Bishop, A. W. Cross, B. R. Johnson, and J. M. Gambetta, Quantum computing with Qiskit, [arXiv:2405.08810](#).
- [31] S. K. Lam, A. Pitrou, and S. Seibert, Numba: A LLVM-based Python JIT compiler, in *Proceedings of the Second Workshop on the LLVM Compiler Infrastructure in HPC* (Association for Computing Machinery, New York, NY, USA, 2015), pp. 1–6.
- [32] J. Lee, W. J. Huggins, M. Head-Gordon, and K. B. Whaley, Generalized unitary coupled cluster wave functions for quantum computation, *J. Chem. Theory Comput.* **15**, 311 (2019).
- [33] P. Jordan and E. Wigner, Über das Paulische Äquivalenzverbot, *Z. Phys.* **51**, 521 (1928).
- [34] Z. Jiang, K. J. Sung, K. Kechedzhi, V. N. Smelyanskiy, and S. Boixo, Quantum algorithms to simulate many-body physics of correlated fermions, *Phys. Rev. Appl.* **9**, 044036 (2018).
- [35] A. Zhao and A. Miyake, Group-theoretic error mitigation enabled by classical shadows and symmetries, *npj Quantum Inf.* **10**, 57 (2024).
- [36] G. van Kempen and L. van Vliet, Mean and variance of ratio estimators used in fluorescence ratio imaging, *Cytometry* **39**, 300 (2000).
- [37] W. J. Huggins, S. McArdle, T. E. O’Brien, J. Lee, N. C. Rubin, S. Boixo, K. B. Whaley, R. Babbush, and J. R. McClean, Virtual distillation for quantum error mitigation, *Phys. Rev. X* **11**, 041036 (2021).
- [38] M. M. Wolf, G. Ortiz, F. Verstraete, and J. I. Cirac, Quantum phase transitions in matrix product systems, *Phys. Rev. Lett.* **97**, 110403 (2006).
- [39] A. Smith, B. Jobst, A. G. Green, and F. Pollmann, Crossing a topological phase transition with a quantum computer, *Phys. Rev. Res.* **4**, L022020 (2022).
- [40] 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).
- [41] K. Klymko, C. Mejuto-Zaera, S. J. Cotton, F. Wudarski, M. Urbanek, D. Hait, M. Head-Gordon, K. B. Whaley, J. Moussa, N. Wiebe, W. A. de Jong, and N. M. Tubman, Real-time evolution for ultracompact hamiltonian eigenstates on quantum hardware, *PRX Quantum* **3**, 020323 (2022).
- [42] W. Kirby, M. Motta, and A. Mezzacapo, Exact and efficient Lanczos method on a quantum computer, *Quantum* **7**, 1018 (2023).
- [43] M. J. D. Powell, A direct search optimization method that models the objective and constraint functions by linear interpolation, in *Advances in Optimization and Numerical Analysis*, edited by S. Gomez and J.-P. Hennart (Springer Netherlands, Dordrecht, 1994), pp. 51–67.
- [44] E. Van Den Berg, Z. K. Mineev, A. Kandala, and K. Temme, Probabilistic error cancellation with sparse Pauli–Lindblad models on noisy quantum processors, *Nat. Phys.* **19**, 1116 (2023).
- [45] Y. Kim, A. Eddins, S. Anand, K. X. Wei, E. Van Den Berg, S. Rosenblatt, H. Nayfeh, Y. Wu, M. Zaletel, K. Temme *et al.*, Evidence for the utility of quantum computing before fault tolerance, *Nature (London)* **618**, 500 (2023).
- [46] L. E. Fischer, T. Dao, I. Tavernelli, and F. Tacchino, Dual-frame optimization for informationally complete quantum measurements, *Phys. Rev. A* **109**, 062415 (2024).
- [47] J. Malmi, K. Korhonen, D. Cavalcanti, and G. García-Pérez, Enhanced observable estimation through classical optimization of informationally overcomplete measurement data: Beyond classical shadows, *Phys. Rev. A* **109**, 062412 (2024).
- [48] A. Caprotti, J. Morris, and B. Dakić, Optimizing quantum tomography via shadow inversion, *Phys. Rev. Res.* **6**, 033301 (2024).
- [49] C. Bertoni, J. Haferkamp, M. Hinsche, M. Ioannou, J. Eisert, and H. Pashayan, Shallow shadows: Expectation estimation using low-depth random Clifford circuits, *Phys. Rev. Lett.* **133**, 020602 (2024).
- [50] H.-Y. Hu, A. Gu, S. Majumder, H. Ren, Y. Zhang, D. S. Wang, Y.-Z. You, Z. Mineev, S. F. Yelin, and A. Seif, Demonstration of robust and efficient quantum property learning with shallow shadows, *Nat. Commun.* **16**, 2943 (2025).
- [51] A. Zhao, N. C. Rubin, and A. Miyake, Fermionic partial tomography via classical shadows, *Phys. Rev. Lett.* **127**, 110504 (2021).
- [52] K. Wan, W. J. Huggins, J. Lee, and R. Babbush, Matchgate shadows for fermionic quantum simulation, *Commun. Math. Phys.* **404**, 629 (2023).
- [53] L. E. Fischer, B. Daniel, T. Dao, I. Tavernelli, and F. Tacchino Raw data for “Large-scale implementation of quantum subspace expansion with classical shadows,” (2026), [10.5281/zenodo.20323161](#).