

Quantum finite-temperature Lanczos method

Gian Gentinetta^{1,2,*}, Friederike Metz^{1,2}, William Kirby³, and Giuseppe Carleo^{1,2}

¹*Institute of Physics, École Polytechnique Fédérale de Lausanne (EPFL), CH-1015 Lausanne, Switzerland*

²*Center for Quantum Science and Engineering, EPFL, Lausanne, Switzerland*

³*IBM Quantum, IBM T. J. Watson Research Center, Yorktown Heights, New York 10598, USA*



(Received 31 March 2026; accepted 19 May 2026; published 11 June 2026)

The computation of thermal properties of quantum many-body systems is a central challenge in our understanding of quantum mechanics. We introduce the quantum finite-temperature Lanczos method (QFTLM), which extends the finite-temperature Lanczos method to quantum computers by combining real-time quantum Krylov methods with efficient preparation of typical states for trace estimation. This approach enables the computation of thermal expectation values while avoiding the exponential scaling inherent to classical exact simulation techniques. Numerical experiments on the transverse-field Ising model show that QFTLM can reproduce thermal observables over a wide temperature range. We further analyze the influence of Krylov dimension, number of trace-estimator states, and Trotter error, and show that suitable regularization is essential for robustness in noisy settings. These results establish QFTLM as a promising framework for finite-temperature quantum simulation.

DOI: [10.1103/f3q5-cm1n](https://doi.org/10.1103/f3q5-cm1n)

I. INTRODUCTION

Understanding properties of quantum many-body systems at finite temperature is fundamental in advancing our knowledge of condensed-matter physics and quantum chemistry [1–3]. To effectively describe experimental phenomena, thermal observables such as magnetization, specific heat, and correlation functions are essential. However, performing such calculations remains a central challenge in modern computational quantum physics due to the Hilbert space dimension growing exponentially with the system size. The finite-temperature regime is especially challenging because observables depend on the full spectrum of the Hamiltonian, rendering many computational tricks employed in ground-state search ineffective in this setting [4,5].

Nevertheless, classical simulation techniques such as quantum Monte Carlo methods and tensor networks have made significant progress toward accessing finite-temperature properties [6]. In particular, the finite-temperature Lanczos method (FTLM) has emerged as a promising candidate [7]. FTLM combines ideas of trace estimation [8–14], where distributions of *typical states* are used to estimate thermal expectation values, with Krylov methods that perform exact diagonalization within a relevant subspace of the exponentially large Hilbert space. While the diagonalization happens in the Krylov subspace, FTLM still requires the preparation of the Krylov basis states in the full Hilbert space. To scale to larger system sizes, classical FTLM thus has to be combined with additional

simulation methods, such as tensor networks, that allow for an efficient representation of many-body quantum states [6].

Quantum computers offer a natural framework for the simulation of many-body quantum systems as the number of qubits required to capture the Hilbert space scales polynomially in the number of particles, avoiding the exponential scaling encountered in exact classical simulation methods [15–20]. Furthermore, there has been remarkable progress in using quantum computers to construct Krylov subspaces for the computation of ground- and excited-state problems [21–55]. A recent work connecting quantum Krylov methods to Szegő quadrature rules describes how thermal observables can be expressed in a Krylov subspace generated through real-time dynamics [56].

In this work, we combine the framework presented in Ref. [56] with the efficient preparation of trace-estimator states to define the quantum finite-temperature Lanczos method (QFTLM) algorithm. We show that this method allows us to compute thermal properties over a wide temperature range. In numerical experiments, we analyze how the algorithm scales with system size, Krylov dimension, and number of trace-estimator states as well as how Trotterization of the dynamics affects the final result.

The remainder of this paper is organized as follows. In Sec. II, we review the classical FTLM and discuss how we extend it to the quantum case. We then test our method on the transverse-field Ising model in Sec. III, where we also investigate the effect of noise on our results.

II. METHODS

A. Finite-temperature Lanczos method

The Lanczos method is a widely used technique to obtain the ground-state and low-lying excited states of sparse Hamiltonians [57]. Extending this algorithm to compute thermal properties of quantum systems is nontrivial, as those generally

*Contact author: gian.gentinetta@epfl.ch

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.

depend on the full spectrum of the Hamiltonian, which cannot be captured effectively in a single Krylov subspace. The finite-temperature Lanczos method addresses this challenge by exploiting properties of *typical states* [7,14,58], known in the classical numerical linear algebra literature as *trace estimators*. Typical states $|\psi_k\rangle$ can be used to estimate expectation values of thermal properties [59], i.e.,

$$\frac{\text{Tr}[e^{-\beta H} O]}{\text{Tr}[e^{-\beta H}]} \approx \frac{\sum_{k=1}^K \langle \psi_k | e^{-\beta H} O | \psi_k \rangle}{\sum_{k=1}^K \langle \psi_k | e^{-\beta H} | \psi_k \rangle}. \quad (1)$$

In fact, there is numerical evidence that for large systems with dense spectra at high temperatures, such observables can be accurately estimated using a single random state [60]. Building on this, the FTLM constructs K distinct Krylov subspaces starting from typical states $|\psi_k\rangle$, in which the projected Hamiltonian is then diagonalized to obtain the projected thermal state.

B. Quantum finite-temperature Lanczos method

In this subsection, we discuss how quantum computers can be employed to perform the FTLM. We discuss how to efficiently prepare typical states and evolve them in time to generate a real-time Krylov basis. We then describe the measurements required to project the thermal state and observables into the subspace.

Ideally, we would like to prepare states drawn randomly from the Haar distribution over the unitary group on n qubits [61]. However, preparing truly random states on a quantum device requires circuit depth exponential in the number of qubits [62,63]. Fortunately, there are several approaches to efficiently approximating the Haar distribution up to a moment T , such that for any quantum state $|\psi\rangle$,

$$\mathbb{E}_{U \sim p_T}[(U|\psi\rangle\langle\psi|U^\dagger)^{\otimes T}] = \mathbb{E}_{U \sim \text{Haar}}[(U|\psi\rangle\langle\psi|U^\dagger)^{\otimes T}], \quad (2)$$

where the distribution p_T is usually referred to as a T -design [64–67]. A particularly efficient approach is described in Ref. [62], where random Clifford circuits are patched together to implement a T -design in $O(\log(n))$ depth.

Instead of preparing fully random states, it is shown in Ref. [61] that the trace estimator converges also for (k -designs of) random diagonal states

$$|\psi_{\text{Q-Hutch}}\rangle = \frac{1}{\sqrt{2^n}} \sum_{j=1}^{2^n} e^{-i\phi_j} |j\rangle, \quad (3)$$

with random complex phases $\phi_j \sim \text{Uniform}(0, 2\pi)$. These states are known as *quantum Hutchinson states* [61] and extend the classical Hutchinson states, where the random phases are real [14]. A 3-design of this distribution can be implemented in $O(n)$ depth using time-evolution circuits of spin glass Hamiltonians [61]. The combination of this simple circuit construction and convergence guarantees stating that the trace estimator scales with $O(1/\sqrt{K})$ motivates our use of these estimators.

In the next step, we construct a Krylov subspace for each $|\psi_k\rangle$, in which we then compute the expectation values in the numerator and denominator of Eq. (1). In classical Krylov methods, the subspace is typically created by repeatedly applying H to the initial quantum state. While this operation is

possible in principle on a quantum computer, it requires fault-tolerant techniques such as qubitization [68]. Instead, we use the quantum Krylov algorithm for Szegő quadrature described in Ref. [56]. The algorithm consists of two main steps. First, a real-time Krylov basis is constructed by evolving $|\psi_k^0\rangle := |\psi_k\rangle$ forward and backward in time using Trotterization of $e^{-i\Delta t H}$. We denote the $D = 2d + 1$ nonorthogonal basis states as

$$|\psi_k^j\rangle = \exp(-i\Delta t H)^j |\psi_k^0\rangle, \quad j \in \{-d, -d+1, \dots, d\}, \quad (4)$$

where the time step Δt is chosen such that $E_j \Delta t \in [-\pi, \pi)$ for all energy eigenvalues E_j , i.e., $\Delta t \leq \pi/||H||$.

We then measure the complex Gram matrix

$$S_{jl} = \langle \psi^j | \psi^l \rangle = \langle \psi^0 | \exp(-i\Delta t H)^{l-j} | \psi^0 \rangle, \quad (5)$$

where we dropped the subscript k on $|\psi_k^j\rangle$ for readability using the Hadamard test [69,70]. The Toeplitz form of S reduces the estimation of the D^2 matrix elements to obtaining one representative element for each of the $D - 1$ diagonals (the main diagonal is fixed to 1 as all basis states are normalized). Noting that $U_{jl} = \langle \psi^j | \exp(-i\Delta t H) | \psi^l \rangle = S_{j+1,l}$, we can directly obtain the projected time-evolution unitary, where the only additional matrix element to be estimated is $U_{-d,d}$.

The Gram matrix and the projected unitary are regularized in postprocessing and the Krylov basis is orthonormalized using S (see Sec. II C) to obtain the unitary $\tilde{U}$ in the orthonormalized Krylov subspace. We can then diagonalize $\tilde{U}$ to obtain the eigenvectors V and eigenvalues $\Lambda = \text{diag}(\lambda_1, \dots, \lambda_D)$ of the time-evolution unitary in the Krylov subspace. Noting that $\lambda_j = e^{-iE_j \Delta t}$, we can directly read out the energy eigenvalues E_j and reconstruct the thermal state in the Krylov subspace. In matrix form, the expectation value is then given as

$$\langle \psi^0 | e^{-\beta H} | \psi^0 \rangle \approx [S^{1/2} V \Lambda^{-\beta/\Delta t} V^\dagger S^{1/2}]_{0,0}. \quad (6)$$

A detailed description of the algorithm is provided in Ref. [56].

To estimate $\langle \psi^0 | e^{-\beta H} \hat{O} | \psi^0 \rangle$ for an observable $\hat{O}$ that is not diagonal in the energy eigenbasis, we also have to measure $\langle \psi^j | \hat{O} | \psi^0 \rangle$ on the quantum device to project the observable into the Krylov subspace. Noting that the projector into the Krylov basis can be written as $\Pi = \sum_{i,j} S_{ij}^{-1} |\psi^i\rangle\langle\psi^j|$, we find

$$\langle \psi^0 | e^{-\beta H} \hat{O} | \psi^0 \rangle \approx \langle \psi^0 | e^{-\beta H} \Pi \hat{O} | \psi^0 \rangle \quad (7)$$

$$= [S^{1/2} V \Lambda^{-\beta/\Delta t} V^\dagger S^{-1/2} \hat{O}]_{0,0}. \quad (8)$$

Among quantum Krylov methods, this algorithm is particularly useful in our case, as it admits convergence guarantees for approximating Gibbs states. In particular, it is shown in Ref. [56] that the error in approximating $\langle \psi_k | e^{-\beta H} | \psi_k \rangle$ is exponentially suppressed with the subspace dimension.

C. Regularization

In practice, the Gram matrix S is often ill-conditioned, leading to numerical instabilities. This becomes an even greater problem if the entries S_{ij} and U_{ij} are measured under noisy conditions. Even under perfect conditions, the matrix U constructed by measuring $U_{ij} = \langle \psi^i | \exp(-i\Delta t H) | \psi^j \rangle$ is not *a priori* unitary. For these reasons, we must regularize the

ALGORITHM 1. Regularized diagonalization.

```

1: input noisy quantum Krylov matrices ( $U, S$ )
2: Choose  $\eta > 0$             $\triangleright$  Tikhonov regularization
3: Choose  $\mu > 0$            $\triangleright$  Singular value truncation
4:  $S \leftarrow S + \eta I$ 
5:  $\tilde{U}' \leftarrow S^{-1/2} U S^{-1/2}$     $\triangleright$  Orthonormalization
6:  $P' \Sigma Q^\dagger = \tilde{U}'$             $\triangleright$  Singular value decomposition
7:  $P \leftarrow P'[:, \Sigma > \mu]$     $\triangleright$  Singular value truncation
8:  $\tilde{U}'' \leftarrow P^\dagger \tilde{U}' P$         $\triangleright$  Transform into truncated basis
9:  $WDR^\dagger = \tilde{U}''$               $\triangleright$  Singular value decomposition
10:  $\tilde{U} \leftarrow WR^\dagger$             $\triangleright$  Project to nearest unitary
11:  $\Lambda = \text{eigvals}(\tilde{U})$           $\triangleright$  Diagonalize
12:  $V' = \text{eigvecs}(\tilde{U})$ 
13:  $V \leftarrow PV'$               $\triangleright$  Transform back into the full basis
14: return  $\Lambda, V$ 

```

measured matrices and orthonormalize the Krylov basis before proceeding with the QFTLM algorithm. A regularization scheme was proposed in Ref. [56], but we find that in practice it is useful to include an additional step in the regularization, which we now describe. Pseudocode is given in Algorithm 1.

As described in Ref. [56], we start by orthonormalizing the Krylov subspace. First, the Gram matrix is regularized with a Tikhonov shift η , ensuring that its eigenvalues are at least of size η . In practice, it is advised to choose η in the order of magnitude of the expected noise on the measured overlaps. We can then safely compute the inverse square root of S and obtain the representation of U in an orthonormal basis for the Krylov space [71]:

$$\tilde{U}' = S^{-1/2} U S^{-1/2}. \quad (9)$$

At this point, our regularization procedure deviates from the analysis in Ref. [56]. As discussed in Ref. [56], in the noiseless case all but one of the eigenvalues of the resulting matrix $\tilde{U}'$ should now have unit length. In practice, however, we observed that noisy inputs often lead to multiple smaller eigenvalues. In Ref. [56], $\tilde{U}'$ is simply projected to the nearest unitary operator (which is unique), but we find that an additional step improves performance in practice. Prior to projecting to the nearest unitary, we prune singular values of $\tilde{U}'$ that lie below a threshold parameter μ . This is achieved by performing a singular value decomposition $\tilde{U}' = P \Sigma Q^\dagger$ and transforming $\tilde{U}'$ with the truncated (and hence nonsquare) transformation P ,

$$\tilde{U}''_{ij} = \sum_{\substack{k,l \\ \text{s.t. } \sigma_k, \sigma_l \geq \mu}} P_{ik}^\dagger \tilde{U}'_{kl} P_{lj}. \quad (10)$$

The motivation for this is that with the exception of the one singular value that is “supposed” to deviate below 1, the deviations of additional singular values indicate that the corresponding dimensions are corrupted by noise, with the degree of the deviation serving as a proxy the degree of the corruption. In our numerics, the optimal threshold μ is typically found to be close to 1. While this additional step does mean that the resulting quadrature rule is no longer formally a Szegő quadrature rule as in Ref. [56], we find that it significantly improves performance in practice.

After this step, we can proceed with the algorithm outlined in Ref. [56] by projecting $\tilde{U}''$ onto the nearest unitary. To this end, we perform another singular value decomposition $\tilde{U}'' = WDR^\dagger$ and set all singular values to unity, yielding $\tilde{U} = WR^\dagger$. Finally, we note that the eigenvectors obtained through diagonalizing $\tilde{U} = V' \Lambda V$ have to be transformed back into the original basis $V = PV'$.

We note that the optimal regularization parameters η and μ vary significantly depending on the input noise. In the results section, we analyze this in more detail and discuss approaches to choose these hyperparameters effectively.

III. RESULTS

We use classical simulations to study the performance of the QFTLM on the transverse-field Ising model (TFIM) given by the Hamiltonian

$$H = \sum_{i=1}^L X_i - \sum_{i=1}^{L-1} Z_i Z_{i+1}, \quad (11)$$




FIG. 1. Thermal expectation values of the energy [panel (a)] and heat capacity [panel (b)] as a function of the temperature for $L = 14$ spins in the TFIM [see Eq. (11)]. We compare the results obtained with QFTLM for varying Krylov dimensions D with the exact solution (in black). The real-time Krylov basis is approximated with four first-order Trotter steps and the expectation values are computed by averaging over $K = 40$ quantum Hutchinson states. In the low-temperature limit, choosing a smaller Krylov subspace leads to significant deviations from the exact solutions. For high temperatures, all simulations qualitatively align with the ground truth.




FIG. 2. Propagation of Trotter error through the QFTLM protocol. In panel (a), we plot the error of the thermal energies per spin as a function of temperature for QFTLM simulations with varying number of Trotter steps used to build the real-time Krylov basis. The Krylov dimension is fixed to $D = 120$ and we average over $K = 40$ quantum Hutchinson states. At low temperatures, the Trotter error is clearly noticeable, whereas at high temperatures other sources dominate the overall error on the thermal energy. In panel (b), we explicitly plot the scaling of the energy error per spin as a function of the number of Trotter steps in the low-temperature limit ($T = 0.025$, indicated by the gray dashed line in the left plot). For a large enough Krylov dimension D , the error scales roughly as $O(\delta t^2)$ (indicated by the black dotted line), where δt denotes the Trotter time step.

where X_i, Z_i are the Pauli operators acting on qubit i . In this section, we investigate how the different sources of error (finite Krylov dimension, trace estimation, Trotter approximation, and shot noise) influence the accuracy of the simulation and how the error scales with the system size L .

In Fig. 1, we show the thermal expectation values of the energy $\langle E \rangle = \text{Tr}[H e^{-\beta H}] / \text{Tr}[e^{-\beta H}]$ [panel (a)] and the heat capacity $\langle C_v \rangle = \beta^2 (\text{Tr}[H^2 e^{-\beta H}] / \text{Tr}[e^{-\beta H}] - \langle E \rangle^2)$ [panel (b)] as functions of the temperature for $L = 14$ spins and varying Krylov dimensions. To build the real-time Krylov basis, we evolve a quantum Hutchinson state in time by approximating $\exp(-i\Delta t H)$ with four first-order Trotter steps [72–74] and repeating this circuit j times to approximate j time steps $\exp(-i\Delta t H)^j$. We have found that choosing $\Delta t = \frac{3\pi}{4|E_0|}$ yields the most accurate results and, hence, use that value for the time step throughout this section. However, we note that if the ground-state energy E_0 is unknown, Δt has to be chosen heuristically to ensure that $\Delta t \leq \pi / \|H\|$.

We iterate over $K = 40$ trace-estimator states to compute the thermal averages. At low temperatures ($T < 1$), the results deviate qualitatively from the exact solution for smaller Krylov dimensions, whereas at high temperatures all simulations overlap with the exact solution. This supports the intuition that at high temperatures, the features in the Gibbs state are smoothed and therefore become easier to approximate at low Krylov dimension.

In Fig. 2, we further analyze the effect of Trotterization in the QFTLM protocol by showing the error of the thermal energy expectation values as a function of the temperature for one, two, and four first-order Trotter steps used to approximate $\exp(-i\Delta t H)$. For now, the Krylov dimension is fixed at $D = 120$ and we again average over $K = 40$ quantum Hutchinson states. As with the Krylov dimension, the error due to choosing large Trotter steps is most apparent in the low-temperature regime. To demonstrate that further, we consider a slice at low temperature [$T = 0.025$, indicated by the

dash-dotted line in panel (a)] and separately display the error as a function of the number of Trotter steps N_T in panel (b). Here, we again include simulations for varying Krylov dimensions and note that for sufficiently large D , the error in the thermal energy roughly scales as $O(1/N_T^2)$, which corresponds to the expected error for a first-order Trotter approximation.

Finally, we investigate how the performance of our method scales with the system size L in the different temperature regimes. Figure 3 shows the error of the thermal energy expectation as a function of L for a varying number of quantum Hutchinson states K used to estimate the trace. The Krylov dimension is kept constant at $D = 20$ over all experiments so that the error due to the finite Krylov dimension is noticeable even for small L and we can study how it scales with the system size. In the low-temperature regime [$T = 0.01$, panel (a)], we observe that the error increases with L . This is to be expected as for $T \rightarrow 0$ the problem essentially reduces to finding the ground state, where the error depends on the gap between the ground-state energy and the first excited state as well as on the overlap between the initial state and the ground state [32,51]. Both of those properties decrease with the system size, and for typical states, which we employ as the initial states in our work, the overlap with the ground state decays exponentially. However, this effect diminishes as the temperature increases and at $T = 100$ [panel (c)] the error in fact decreases with L . This provides evidence that, like FTLM [59], QFTLM works particularly well at high temperatures and large system sizes. In practice, the overlaps computed on the quantum device can only be estimated to finite precision. Even on a fault-tolerant quantum computer, those quantities will be affected by finite sampling noise. To understand how the QFTLM algorithm performs under noisy conditions, we add Gaussian noise to the Gram matrix S and projected unitary U and observe how this error propagates through the algorithm. To this end, we exploited the Toeplitz form of S to construct S and U with a total of D noisy overlaps. This not




FIG. 3. Scaling of the QFTLM error with system size L . We plot the error in the energy expectation values per spin at three temperatures, $T = 0.01$, 1 , and 100 , in panels (a), (b), and (c), respectively, as a function of the number of spins L , for a varying number of quantum Hutchinson states K . The Krylov subspace dimension is fixed to $D = 20$ and the number of Trotter steps to 4 . At low temperatures, the error at fixed Krylov dimension increases with the system size. Conversely, at high temperatures, we observe a decrease in the error as the number of spins increases.

only reduces the complexity of the estimation but also directly enforces Hermiticity of S and fixes its diagonal entries to 1 . The Gaussian noise acts as a proxy for shot noise because

the binomial distributions arising from measuring the real and imaginary parts of the complex overlap are well approximated by Gaussian distributions in the large-shot limit. By the central limit theorem, the corresponding estimator fluctuations scale as $O(1/\sqrt{R})$, where R is the number of measurement shots.

We again look at $L = 14$ spins in the TFIM described in Eq. (11) for a Krylov dimension of $D = 120$ and average over $K = 40$ quantum Hutchinson states. In Fig. 4, we show the error per spin of the thermal energy expectation values as a function of the temperature for varying noise levels. We show results for a simulation where the time evolution is performed exactly [panel (a)] as well as one where we approximate the $e^{-iH\Delta t}$ using four Trotter steps [panel (b)]. In the exact case, we observe that the noise is the dominant source of error at low temperatures. At high temperatures ($T > 1$), Gaussian noise with standard deviation of 10^{-6} is small enough that overall it is not the dominant source of error. In the simulation using Trotterized time evolution, the error for noise levels below 10^{-6} is comparable throughout the whole temperature range since the Trotterization is the dominant source of error in this case. Still, we note that for standard deviations above $\sigma = 10^{-5}$, the simulation deviates significantly from the exact solution. On an ideal quantum computer, we expect to require $O(1/\sigma^2)$ shots in order to measure the overlaps with precision σ , which implies a significant sampling overhead for the QFTLM algorithm.

For simulations affected by shot noise, the regularization procedure discussed in Sec. II C becomes particularly important. While throughout our experiments, we observed that choosing the singular-value truncation threshold in Eq. (10) around $\mu \approx 0.9$ yields the best results, the ideal choice of the Tikhonov shift η strongly depends on the noise level. In Fig. 5, we show the mean error on the thermal energies as a function of η . We observe that the optimal result is achieved when η is chosen two orders of magnitude larger than the standard deviation of the Gaussian noise applied to the measured overlaps. If the shift is chosen too small, the computation of $S^{-1/2}$ fails, yielding huge errors in the observed energies. On the other hand, we note that it is acceptable to choose η to be larger than the optimal value as the error does not significantly increase until $\eta \gtrsim 10^{-4}$. Ideally, η should thus be chosen generously above 100 times the expected noise level. It is important to note that the choice of regularization happens in postprocessing and hence can be optimized given the noisy data.

In particular, in cases where the ground-state energy is known (or can be computed using other computational methods), η and μ can be fine-tuned to minimize the error on the ground state obtained in the diagonalization of $\tilde{U}$. We note that obtaining the ground-state energy of a system is often tractable using classical numerical techniques even in cases where thermal properties are classically challenging [75,76]. Such cases can be well adapted for the above tuning scheme. Even if the error for a ground or very low-temperature state cannot be made small, the primary goal in this tuning is to avoid the ill-conditioned regime, which can be assessed using a classical baseline.

IV. CONCLUSION

We introduce the QFTLM, combining the classical FTLM framework with real-time quantum Krylov methods and




FIG. 4. Effect of noise on the error of thermal energies obtained through QFTLM. Gaussian noise with varying standard deviation is added to the measured Gram matrix S and projected unitary U . In panel (a), the time evolution is performed exactly without Trotterization. In this case, the noise is the dominant source of error at low temperatures, whereas for sufficiently small standard deviations the error saturates at high temperatures. In panel (b), we perform the time evolution by splitting $e^{-iH\Delta t}$ into four Trotter steps. Now, the Trotter error dominates at low temperatures for small enough noise levels. Colorbar ticks indicate the simulated noise levels.

efficient trace estimation using typical states. By averaging expectation values over multiple Krylov iterations starting from independent quantum states, we are able to compute thermal properties, effectively extending the Krylov framework beyond low-lying excited states. By employing a quantum computer to prepare the Krylov basis and measure the Gram matrix, we circumvent the curse of dimensionality encountered in classical simulation techniques.

We demonstrate through numerical experiments on the transverse-field Ising model that QFTLM can accurately reproduce thermal expectation values over a wide temperature range. Our results highlight the complementary roles of the

Krylov dimension, the number of trace-estimator states, and the accuracy of the time-evolution operator. We observe that the magnitude of the Trotterization error and Krylov dimension particularly matter in the low-temperature regime, whereas at high temperatures typicality becomes increasingly important. Of particular interest is the scaling of the error with the system size, where we find that at high temperatures the error per spin at fixed Krylov dimension decreases as the number of spins increases. This suggests a beneficial scaling for the QFTLM algorithm once it is extended past the size of exact diagonalization.

Furthermore, we investigate how noise in the quantum measurement propagates through the QFTLM algorithm. To this end, we design a regularization procedure to separate signal from noise. We observe that the QFTLM is robust to small noise levels, but it is clear that a large number of shots will be required to resolve the measured overlaps to a precision that allows for accurate results.

In summary, QFTLM provides a framework for computing finite-temperature properties of quantum many-body systems on quantum computers. The method is applicable to arbitrary quantum Hamiltonians beyond the TFIM studied in this paper. We expect QFTLM to serve as a useful tool in estimating thermal expectations in condensed-matter physics and quantum chemistry on near- to mid-term quantum devices.




FIG. 5. Importance of appropriate regularization for noisy simulations. We plot the absolute energy error averaged over all temperatures (between 10^{-2} and 10^2) as a function of the Tikhonov shift η chosen to regularize the Gram matrix (line 4 in Algorithm 1). We find that for each noise level there is a threshold (roughly a factor of 100 above the noise level) defining the minimal η above which we obtain reasonable errors in the energy. Below that threshold, the error spikes up to $O(1)$, whereas above the threshold the error remains low up to $\eta \gtrsim 10^{-4}$. Colorbar ticks indicate the simulated noise levels.

ACKNOWLEDGMENT

This research was supported by the NCCR MARVEL, a National Centre of Competence in Research, funded by the Swiss National Science Foundation (Grant No. 205602).

DATA AVAILABILITY

The code and data required to reproduce the figures shown in this work are available at Ref. [77].

- [1] N. D. Mermin, Thermal properties of the inhomogeneous electron gas, *Phys. Rev.* **137**, A1441 (1965).
- [2] A. Georges, G. Kotliar, W. Krauth, and M. J. Rozenberg, Dynamical mean-field theory of strongly correlated fermion systems and the limit of infinite dimensions, *Rev. Mod. Phys.* **68**, 13 (1996).
- [3] A. F. White and G. K.-L. Chan, Finite-temperature coupled cluster: Efficient implementation and application to prototypical systems, *J. Chem. Phys.* **152**, 224104 (2020).
- [4] G. Harsha, T. M. Henderson, and G. E. Scuseria, Thermofield theory for finite-temperature quantum chemistry, *J. Chem. Phys.* **150**, 154109 (2019).
- [5] C. Sun, U. Ray, Z.-H. Cui, M. Stoudenmire, M. Ferrero, and G. K.-L. Chan, Finite-temperature density matrix embedding theory, *Phys. Rev. B* **101**, 075131 (2020).
- [6] A. Sinha, M. M. Rams, and J. Dziarmaga, Efficient representation of minimally entangled typical thermal states in two dimensions via projected entangled pair states, *Phys. Rev. B* **109**, 045136 (2024).
- [7] J. Jaklič and P. Prelovšek, Lanczos method for the calculation of finite-temperature quantities in correlated systems, *Phys. Rev. B* **49**, 5065 (1994).
- [8] G. H. Golub and G. Meurant, *Matrices, Moments and Quadrature with Applications*, Princeton Series in Applied Mathematics (Princeton University Press, Princeton, NJ, 2010).
- [9] H. Avron and S. Toledo, Randomized algorithms for estimating the trace of an implicit symmetric positive semi-definite matrix, *J. ACM* **58**, 1 (2011).
- [10] G. H. Golub and C. F. Van Loan, *Matrix Computations*, 4th ed. (Johns Hopkins University Press, Philadelphia, PA, 2013).
- [11] F. Roosta-Khorasani and U. Ascher, Improved bounds on sample size for implicit matrix trace estimators, *Found. Comput. Math.* **15**, 1187 (2015).
- [12] S. Ubaru, J. Chen, and Y. Saad, Fast estimation of $\text{tr}(f(A))$ via stochastic Lanczos quadrature, *SIAM J. Matrix Anal. Appl.* **38**, 1075 (2017).
- [13] P.-G. Martinsson and J. A. Tropp, Randomized numerical linear algebra: Foundations and algorithms, *Acta Numer.* **29**, 403 (2020).
- [14] M. Hutchinson, A stochastic estimator of the trace of the influence matrix for Laplacian smoothing splines, *Commun. Stat. Simul. Comput.* **18**, 1059 (1989).
- [15] R. P. Feynman, Simulating physics with computers, *Int. J. Theor. Phys.* **21**, 467 (1982).
- [16] 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).
- [17] A. Chiesa, F. Tacchino, M. Grossi, P. Santini, I. Tavernelli, D. Gerace, and S. Carretta, Quantum hardware simulating four-dimensional inelastic neutron scattering, *Nat. Phys.* **15**, 455 (2019).
- [18] F. Arute *et al.*, Hartree-Fock on a superconducting qubit quantum computer, *Science* **369**, 1084 (2020).
- [19] F. Arute *et al.*, Observation of separated dynamics of charge and spin in the Fermi-Hubbard model, [arXiv:2010.07965](https://arxiv.org/abs/2010.07965).
- [20] C. Neill *et al.*, Accurately computing the electronic properties of a quantum ring, *Nature (London)* **594**, 508 (2021).
- [21] J. R. McClean, M. E. Kimchi-Schwartz, J. Carter, and W. A. de Jong, Hybrid quantum-classical hierarchy for mitigation of decoherence and determination of excited states, *Phys. Rev. A* **95**, 042308 (2017).
- [22] J. I. Colless, V. V. Ramasesh, D. Dahlen, M. S. Blok, M. E. Kimchi-Schwartz, J. R. McClean, J. Carter, W. A. de Jong, and I. Siddiqi, Computation of molecular spectra on a quantum processor with an error-resilient algorithm, *Phys. Rev. X* **8**, 011021 (2018).
- [23] R. M. Parrish and P. L. McMahon, Quantum filter diagonalization: Quantum eigendecomposition without full quantum phase estimation, [arXiv:1909.08925](https://arxiv.org/abs/1909.08925).
- [24] 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).
- [25] 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).
- [26] W. J. Huggins, J. Lee, U. Baek, B. O'Gorman, and K. B. Whaley, A non-orthogonal variational quantum eigensolver, *New J. Phys.* **22**, 073009 (2020).
- [27] N. H. Stair, R. Huang, and F. A. Evangelista, A multireference quantum Krylov algorithm for strongly correlated electrons, *J. Chem. Theory Comput.* **16**, 2236 (2020).
- [28] 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).
- [29] O. Kyriienko, Quantum inverse iteration algorithm for programmable quantum simulators, *npj Quantum Inf.* **6**, 7 (2020).
- [30] J. Cohn, M. Motta, and R. M. Parrish, Quantum filter diagonalization with compressed double-factorized Hamiltonians, *PRX Quantum* **2**, 040352 (2021).
- [31] N. Yoshioka, H. Hakoshima, Y. Matsuzaki, Y. Tokunaga, Y. Suzuki, and S. Endo, Generalized quantum subspace expansion, *Phys. Rev. Lett.* **129**, 020502 (2022).
- [32] E. N. Epperly, L. Lin, and Y. Nakatsukasa, A theory of quantum subspace diagonalization, *SIAM J. Matrix Anal. Appl.* **43**, 1263 (2022).
- [33] K. Seki and S. Yunoki, Quantum power method by a superposition of time-evolved states, *PRX Quantum* **2**, 010333 (2021).
- [34] T. A. Bespalova and O. Kyriienko, Hamiltonian operator approximation for energy measurement and ground-state preparation, *PRX Quantum* **2**, 030318 (2021).
- [35] T. E. Baker, Lanczos recursion on a quantum computer for the Green's function and ground state, *Phys. Rev. A* **103**, 032404 (2021).
- [36] T. E. Baker, Block Lanczos method for excited states on a quantum computer, *Phys. Rev. A* **110**, 012420 (2024).
- [37] C. L. Cortes and S. K. Gray, Quantum Krylov subspace algorithms for ground- and excited-state energy estimation, *Phys. Rev. A* **105**, 022417 (2022).
- [38] 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).
- [39] F. Jamet, A. Agarwal, and I. Rungger, Quantum subspace expansion algorithm for Green's functions, [arXiv:2205.00094](https://arxiv.org/abs/2205.00094).
- [40] U. Baek, D. Hait, J. Shee, O. Leimkuhler, W. J. Huggins, T. F. Stetina, M. Head-Gordon, and K. B. Whaley, Say

- no to optimization: A nonorthogonal quantum eigensolver, *PRX Quantum* **4**, 030307 (2023).
- [41] N. V. Tkachenko, L. Cincio, A. I. Boldyrev, S. Tretiak, P. A. Dub, and Y. Zhang, Quantum Davidson algorithm for excited states, *Quantum Sci. Technol.* **9**, 035012 (2024).
- [42] G. Lee, D. Lee, and J. Huh, Sampling error analysis in quantum Krylov subspace diagonalization, *Quantum* **8**, 1477 (2024).
- [43] Z. Zhang, A. Wang, X. Xu, and Y. Li, Measurement-efficient quantum Krylov subspace diagonalisation, *Quantum* **8**, 1438 (2024).
- [44] W. Kirby, M. Motta, and A. Mezzacapo, Exact and efficient Lanczos method on a quantum computer, *Quantum* **7**, 1018 (2023).
- [45] Y. Shen, K. Klymko, J. Sud, D. B. Williams-Young, W. A. d. Jong, and N. M. Tubman, Real-time Krylov theory for quantum computing algorithms, *Quantum* **7**, 1066 (2023).
- [46] B. Yang, N. Yoshioka, H. Harada, S. Hakkaku, Y. Tokunaga, H. Hakoshima, K. Yamamoto, and S. Endo, Resource-efficient generalized quantum subspace expansion, *Phys. Rev. Appl.* **23**, 054021 (2025).
- [47] R. Yang, T. Wang, B.-N. Lu, Y. Li, and X. Xu, Shadow-based quantum subspace algorithm for the nuclear shell model, *Phys. Rev. A* **111**, 012620 (2025).
- [48] Y. Ohkura, S. Endo, T. Satoh, R. V. Meter, and N. Yoshioka, Leveraging hardware-control imperfections for error mitigation via generalized quantum subspace, [arXiv:2303.07660](https://arxiv.org/abs/2303.07660).
- [49] 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).
- [50] L. W. Anderson, M. Kiffner, T. O’Leary, J. Crain, and D. Jaksch, Solving lattice gauge theories using the quantum Krylov algorithm and qubitization, *Quantum* **9**, 1669 (2025).
- [51] W. Kirby, Analysis of quantum Krylov algorithms with errors, *Quantum* **8**, 1457 (2024).
- [52] 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).
- [53] A. Byrne, W. Kirby, K. M. Soodhalter, and S. Zhuk, The quantum super-Krylov method, [arXiv:2412.17289](https://arxiv.org/abs/2412.17289).
- [54] J. Yu, J. R. Moreno, J. T. Iosue, L. Bertels, D. Claudino, B. Fuller, P. Groszkowski, T. S. Humble, P. Jurcevic, W. Kirby, T. A. Maier, M. Motta, B. Pokharel, A. Seif, A. Shehata, K. J. Sung, M. C. Tran, V. Tripathi, A. Mezzacapo, and K. Sharma, Quantum-centric algorithm for sample-based Krylov diagonalization, [arXiv:2501.09702](https://arxiv.org/abs/2501.09702).
- [55] S. Piccinelli, A. Baiardi, S. Barison, M. Rossmannek, A. C. Vazquez, F. Tacchino, S. Mensa, E. Altamura, A. Alavi, M. Motta, J. Robledo-Moreno, W. Kirby, K. Sharma, A. Mezzacapo, and I. Tavernelli, Quantum chemistry with provable convergence via randomized sample-based Krylov quantum diagonalization, [arXiv:2508.02578](https://arxiv.org/abs/2508.02578).
- [56] W. Kirby, Y. Shen, D. Camps, A. Chowdhury, K. Klymko, and R. V. Beeumen, Quantum Krylov algorithm for Szegő quadrature, [arXiv:2509.19195](https://arxiv.org/abs/2509.19195).
- [57] C. Lanczos, An iteration method for the solution of the eigenvalue problem of linear differential and integral operators, *J. Res. Natl. Bur. Stand.* **45**, 255 (1950).
- [58] T. Iitaka and T. Ebisuzaki, Random phase vector for calculating the trace of a large matrix, *Phys. Rev. E* **69**, 057701 (2004).
- [59] J. Schnack, J. Richter, and R. Steinigeweg, Accuracy of the finite-temperature Lanczos method compared to simple typicality-based estimates, *Phys. Rev. Res.* **2**, 013186 (2020).
- [60] J. Schnack, J. Schulenburg, and J. Richter, Magnetism of the $N = 42$ kagome lattice antiferromagnet, *Phys. Rev. B* **98**, 094423 (2018).
- [61] Y. Shen, K. Klymko, E. Rabani, N. M. Tubman, D. Camps, R. V. Beeumen, and M. Lindsey, Diagonal state designs with reconfigurable real-time circuits, *Phys. Rev. Res.* **8**, 013042 (2026).
- [62] T. Schuster, J. Haferkamp, and H.-Y. Huang, Random unitaries in extremely low depth, *Science* **389**, 92 (2025).
- [63] S. Aaronson, The complexity of quantum states and transformations: From quantum money to black holes, [arXiv:1607.05256](https://arxiv.org/abs/1607.05256).
- [64] P. Hayden, D. Leung, P. W. Shor, and A. Winter, Randomizing quantum states: Constructions and applications, *Commun. Math. Phys.* **250**, 371 (2004).
- [65] C. Dankert, R. Cleve, J. Emerson, and E. Livine, Exact and approximate unitary 2-designs and their application to fidelity estimation, *Phys. Rev. A* **80**, 012304 (2009).
- [66] D. Gross, K. Audenaert, and J. Eisert, Evenly distributed unitaries: On the structure of unitary designs, *J. Math. Phys.* **48**, 052104 (2007).
- [67] A. W. Harrow, A. Hassidim, and S. Lloyd, Quantum algorithm for linear systems of equations, *Phys. Rev. Lett.* **103**, 150502 (2009).
- [68] G. H. Low and I. L. Chuang, Hamiltonian simulation by qubitization, *Quantum* **3**, 163 (2019).
- [69] R. Cleve, A. Ekert, C. Macchiavello, and M. Mosca, Quantum algorithms revisited, *Proc. R. Soc. London, Ser. A* **454**, 339 (1998).
- [70] D. Aharonov, V. Jones, and Z. Landau, A polynomial quantum algorithm for approximating the Jones polynomial, in *Proceedings of the 38th Annual ACM Symposium on Theory of Computing* (ACM Press, New York, NY, 2006), pp. 427–436.
- [71] P.-O. Löwdin, On the nonorthogonality problem, *Adv. Quantum Chem.* **5**, 185 (1970).
- [72] N. Hatano and M. Suzuki, Finding exponential product formulas of higher orders, in *Quantum Annealing and Other Optimization Methods*, edited by A. Das and B. K. Chakrabarti (Springer, Berlin, 2005), pp. 37–68.
- [73] S. Lloyd, Universal quantum simulators, *Science* **273**, 1073 (1996).
- [74] A. M. Childs, Y. Su, M. C. Tran, N. Wiebe, and S. Zhu, Theory of Trotter error with commutator scaling, *Phys. Rev. X* **11**, 011020 (2021).
- [75] S. R. White, Density matrix formulation for quantum renormalization groups, *Phys. Rev. Lett.* **69**, 2863 (1992).
- [76] U. Schollwöck, The density-matrix renormalization group in the age of matrix product states, *Ann. Phys.* **326**, 96 (2011).
- [77] G. Gentinetta and F. Metz, *Data for Manuscript on the Quantum Finite Temperature Lanczos Method* (Materials Cloud Archive, 2026), <https://github.com/gentinettatagian/QuantumFiniteTemperatureLanczos>.