



Cite this: DOI: 10.1039/d5sc09162d

All publication charges for this article have been paid for by the Royal Society of Chemistry

Constrained shadow tomography for molecular simulation on quantum devices

Irma Avdic,^a Yuchen Wang,^{†a} Michael Rose,^{†a} Lillian I. Payne Torres,^a Anna O. Schouten,^a Kevin J. Sung^b and David A. Mazziotti^{*,a}

Quantum state tomography is a fundamental task in quantum information science, enabling detailed characterization of correlations, entanglement, and electronic structure in quantum systems. However, its exponential measurement and computational demands limit scalability, motivating efficient alternatives such as classical shadows, which enable accurate prediction of many observables from randomized measurements. In this work, we introduce a bi-objective semidefinite programming approach for constrained shadow tomography, designed to reconstruct the two-particle reduced density matrix (2-RDM) from noisy or incomplete shadow data. By integrating N -representability constraints and nuclear-norm regularization into the optimization, the method builds an N -representable 2-RDM that balances fidelity to the shadow measurements with energy minimization. This unified framework mitigates noise and sampling errors while enforcing physical consistency in the reconstructed states. Numerical and hardware results demonstrate that the approach significantly improves accuracy, noise resilience, and scalability, providing a robust foundation for physically consistent fermionic state reconstruction in realistic quantum simulations.

Received 23rd November 2025

Accepted 15th March 2026

DOI: 10.1039/d5sc09162d

rsc.li/chemical-science

1. Introduction

Reconstructing a system's quantum state from measurement data, a process known as quantum state tomography,¹⁻⁴ is foundational to quantum information science, many-body physics, and quantum chemistry. Accurate knowledge of a system's quantum state enables quantitative characterization of correlations, entanglement, and electronic structure properties, all of which are important for performing quantum simulations, benchmarking quantum hardware, and interpreting experimental outcomes. Quantum state tomography provides the most direct route to validating state preparation, certifying entanglement in quantum devices, and probing nonclassical correlations in condensed matter and chemical systems. However, conventional full-state tomography requires measurement and computational resources that scale exponentially with system size, rendering it infeasible for all but the smallest quantum systems.⁵⁻⁸

To overcome this limitation, a wide range of scalable tomographic approaches have been developed. These methods aim to infer only the most relevant or physically meaningful aspects of a quantum state rather than reconstructing the full

wavefunction. Such approaches include compressed sensing,^{9,10} neural-network-based tomography,^{11,12} direct fidelity estimation,¹³⁻¹⁵ and maximum-entropy-inspired methods based on restricted Boltzmann machines,¹⁶ among others. A particularly powerful and widely adopted framework is that of classical shadows,^{17,18} which enables efficient estimation of a large number of observables from randomized measurements. The shadow tomography protocol constructs a compact classical representation, termed a "shadow", from a set of randomized qubit or fermionic measurements, from which expectation values of many observables can be predicted with polynomial or even logarithmic scaling in system size. This paradigm has been successfully applied in a variety of contexts, including entanglement estimation,¹⁹ Hamiltonian learning,²⁰⁻²² and reduced density matrix (RDM) reconstruction for quantum chemistry.^{23,24}

When specialized to fermionic systems, shadow-based methods can leverage rich physical structure, such as anti-symmetry,²⁵ particle-number conservation,²⁴ spin symmetries, and Gaussian unitary transformations,^{23,26,27} to further improve sampling efficiency and reconstruction fidelity. These fermionic shadow protocols enable scalable characterization of electronic structure properties and correlation functions relevant to molecular and solid-state systems, potentially bridging the gap between quantum simulation and classical post-processing. However, despite their theoretical efficiency, existing fermionic shadow tomography approaches face significant practical challenges. In particular, the presence of sampling noise,

^aDepartment of Chemistry, The James Franck Institute, The University of Chicago, Chicago, IL 60637, USA. E-mail: damazz@uchicago.edu

^bIBM Quantum, IBM T.J. Watson Research Center, Yorktown Heights, NY 10598, USA. E-mail: kevinsung@ibm.com

[†] These authors contributed equally to this work.



gate infidelity, and finite measurement statistics on quantum devices can severely degrade reconstruction accuracy.

In this work, we introduce an improved and optimized constrained shadow tomography framework, building on prior work,^{28,29} specifically tailored for the simulation of realistic fermionic systems on quantum devices. Our method formulates the reconstruction of the two-particle reduced density matrix (2-RDM) as a bi-objective semidefinite optimization, extending shadow-based tomography through a constrained formulation designed to improve robustness to noise and limited measurement data. By incorporating physical N -representability constraints directly into the optimization and adding a nuclear-norm regularization to the energy minimization, the method reconstructs a shadow-consistent 2-RDM that satisfies the N -representability conditions through an optimization that balances energy minimization against fidelity to the measured shadow data. This unified semidefinite program simultaneously mitigates errors arising from finite sampling and hardware noise while enforcing physical consistency in the resulting RDMs. Consequently, the framework can be viewed as a noise-aware, state-tailored extension of variational RDM theory,^{30–50} capable of producing physically valid density matrices under realistic quantum measurement conditions.

Through comprehensive numerical analysis, we explore how measurement requirements scale with system size, orbital basis dimension, and the choice of N -representability constraints, identifying practical regimes where constrained tomography yields superior performance over other approaches. Finally, we validate constrained shadow tomography through computations performed on IBM's superconducting quantum device, confirming that the advantages of constraint-enforced shadow tomography persist under realistic experimental conditions involving gate noise, readout errors, and limited sampling depth. Our results highlight the importance of integrating physical priors and optimization-based constraints into quantum state reconstruction frameworks, paving the way for accurate and scalable molecular simulation on quantum hardware.

2. Theory

2.1. Shadow tomography for fermionic systems

Characterizing the full state of a quantum system requires exponentially scaling resources, making it infeasible for practical applications. Tomography schemes based on random measurements have emerged as more computationally accessible alternatives,^{51–53} with classical shadow tomography being among the most widely used.^{17,18} Classical shadow tomography allows for the reconstruction of a quantum state from repeated measurements in different unitary bases, using resources that scale logarithmically with the system size in some cases.²³ By applying random unitary operations, U , to the quantum state, $|\psi\rangle$, followed by a computational basis measurement, $|b\rangle$, a “classical shadow” estimator can be constructed,

$$\hat{\rho}_i = \mathcal{M}^{-1}(U_i^\dagger |b_i\rangle \langle b_i| U_i), \quad (1)$$

where \mathcal{M} is the measurement channel. The estimator can then be used to predict the expectation values of many observables with provable guarantees,

$$\langle O \rangle_{\text{shadow}} = \frac{1}{M} \sum_{i=1}^M \text{Tr}(O \hat{\rho}_i), \quad (2)$$

which converges to $\text{Tr}(O\rho)$ as the number of measurements M increases. For fermionic systems, the choice of measurement unitaries is key to an optimal protocol. For example, particle-number conserving transformations or fermionic Gaussian unitaries (FGUs) preserve antisymmetry and can be implemented efficiently on qubit hardware through linear optical or Majorana representations.^{23,25,27} These structured unitaries significantly reduce circuit depth and improve measurement fidelity compared to fully random Clifford unitaries.

The statistical efficiency of classical shadow tomography is further characterized by the shadow norm, which quantifies the variance of observable estimators under random measurements.¹⁸ For a given observable O , the estimation error scales as

$$\text{Var}[\langle O \rangle_{\text{shadow}}] \leq \frac{\|O\|_{\text{sh}}^2}{M}, \quad (3)$$

where $\|O\|_{\text{sh}}$ is the shadow norm determined by the chosen measurement ensemble. A smaller shadow norm implies that fewer measurement samples are required to achieve a target precision, emphasizing the importance of designing measurement schemes that minimize $\|O\|_{\text{sh}}$ for low-rank or local fermionic operators.

In the context of electronic structure and many-fermion simulations, the goal is often not full state reconstruction but accurate estimation of low-order observables such as the 2-RDM. Classical shadow tomography provides an efficient statistical framework for directly estimating these quantities from measurement data, enabling integration with post-processing schemes and physical consistency.

Since fermionic wavefunctions prohibitively scale exponentially with particle number, their state tomography is often simplified to an estimation of p -particle reduced density matrices (p -RDMs). If represented in the qubit basis, estimating all p -RDMs on n fermionic modes requires $\mathcal{O}((2n)^p/\epsilon^2)$ measurements.^{54,55} Using FGUs, the sample complexity can be reduced to $\mathcal{O}\left(\left(\frac{\eta}{p}\right)p^{3/2}\log(n)/\epsilon^2\right)$.²³ Considering particle-number symmetry, the required number of samples can further be reduced to $\mathcal{O}(\eta^p/\epsilon^2)$,²⁴ with η particles. Here, the estimation error remains independent of the full Hilbert space dimension.

2.2. Constrained shadow tomography of the 2-RDM

In fermionic shadow tomography, the conservation of parity and particle number symmetry ensures compatibility with the physically relevant symmetry sectors of many-body Hamiltonians, leading to reduced sample complexity. Similarly, ensuring the estimated p -RDMs correspond to N -particle wavefunctions, *i.e.*, that they are N -representable,^{30,41,47,56–62} allows for a reduction in



measurement overhead as previously demonstrated by some of the authors.²⁸ Since the electronic Hamiltonian contains at most pairwise interactions, we restrict demonstration of the present work to 2-RDMs. The 2-RDM elements of an N -electron system arise from the integration of the N -particle density matrix over all particles except two and can be expressed as

$${}^2D_{kl}^{ij} = \langle \Psi | \hat{a}_i^\dagger \hat{a}_j^\dagger \hat{a}_l \hat{a}_k | \Psi \rangle, \quad (4)$$

where \hat{a}_i^\dagger creates a particle in orbital i and \hat{a}_i annihilates a particle in orbital i . The classical shadow representation of the 2-RDM can be expressed as

$$S_n^{pq} = \langle \Psi | \hat{U}_n^\dagger \hat{a}_p^\dagger \hat{a}_q^\dagger \hat{a}_q \hat{a}_p \hat{U}_n | \Psi \rangle, \quad (5)$$

where $\hat{U}_n = \exp\left(\sum_{uv} A_n^{uv} \hat{a}_u^\dagger \hat{a}_v\right)$. Here, the indices denote spin

orbitals, A_n is a one-body anti-Hermitian matrix, and n is the shadow index. Each shadow n corresponds to measuring all diagonal elements of the 2-RDM after applying the one-body unitary transformation, \hat{U}_n , to the wave function. The one-body unitaries are generated by random sampling according to the Haar measure,⁶³ effectively rotating the orbitals into a new basis. From eqn (5), the shadow measurements can be expressed in terms of the 2-RDM elements as

$$S_n^{pq} = \sum_{ijkl} U_n^{pi} U_n^{qj} {}^2D_{kl}^{ij} U_n^{ql} U_n^{pk}, \quad (6)$$

where $U_n = \exp(A_n)$. A sufficiently large collection of such shadows defines a system of equations that uniquely determines the 2-RDM. Consequently, this approach can be used to reconstruct the 2-RDM (or any one- or two-body expectation value) of the original quantum state. Importantly, since the measurements only involve the classical components of the 2-RDM, which commute and can therefore be measured simultaneously, the overall cost of the tomography is substantially reduced.

By defining an objective functional of the 2-RDM, ($J[{}^2D]$), such as the nuclear norm, Frobenius norm,⁶⁴ or the energy expectation value,^{39,65,66} a convex optimization problem can be formulated with respect to the shadow constraints,

$$\begin{aligned} & \min_{{}^2D \in N_2\tilde{P}} J[{}^2D] \\ \text{such that } & S_{pq}^n = \left((U \otimes U)^2 D (U \otimes U)^T \right)_{pq}^{pq}, \end{aligned} \quad (7)$$

for $n \in [0, m]$, where $N_2\tilde{P}$ denotes the convex set of approximately N -representable 2-RDMs. For a quantum many-body system with at most pairwise interactions, minimizing the energy functional $E[{}^2D]$ is particularly appealing due to its direct relevance to molecular simulation. The 2-RDM can further be used to estimate the energy gradients^{67,68} and multipole moments.⁶⁹ However, to enable practical implementation on current quantum devices, the formulation must be augmented to properly account for quantum measurement errors.

Motivated by the need to consider two complementary objectives: (1) to obtain an optimal solution to a general functional of the 2-RDM and (2) to minimize measurement errors (e.g., shot/gate noise, readout errors, limited sampling depth,

etc.) affecting the 2-RDM, we cast eqn (7) into the following bi-objective semidefinite-programming optimization problem:^{70,71}

$$\min_{{}^2D \in N_2\tilde{P}} J[{}^2D] + w \| {}^2D - {}^2\tilde{D} \|_* \quad (8)$$

$$\text{subject to } \tilde{S}_{pq}^n = \left((U \otimes U)^2 \tilde{D} (U \otimes U)^T \right)_{pq}^{pq}, \quad (9)$$

where

$${}^2\tilde{D} = {}^2D + \mathbf{E}_1 - \mathbf{E}_2. \quad (10)$$

Here, $w > 0$ is a penalty weight (i.e., regularization parameter), $\| \cdot \|_*$ denotes nuclear norm, and $\mathbf{E}_i \geq 0$ are positive semidefinite slack “error” matrices. The penalty parameter w controls the trade-off between energy minimization and physical deviations of the 2-RDM from the shadow measurements. A smaller w allows for a lower energy value but with potentially larger shadow-constraint violations. A similar relaxation strategy was previously employed by some of the authors to infer time-dependent quantum states from measurement data using matrix completion techniques and a modified Frobenius norm within a multi-objective SDP.⁷² Related approaches have since been investigated on quantum computers using purification of the 2-RDM⁶⁴ and the Quantum Approximate Optimization Algorithm.⁷³

The second objective function in eqn (9) minimizes the nuclear norm, or the sum of the singular values of the slack matrix, over the constraint set including the shadow constraints. The nuclear norm serves as a convex and computationally tractable surrogate for matrix rank, and efficient first-order algorithms based on singular value thresholding have been developed to solve such problems at large scale.⁶⁶ Being convex and easy to optimize, the nuclear norm offers the best convex approximation of the rank function for matrices whose largest singular value is no greater than one.⁷⁴ When the matrix variable is symmetric and positive semidefinite, the method relaxes to the trace heuristic, i.e., $\|X\|_* = \text{Tr}(X)$, frequently used in systems and control studies.^{75,76}

For practical implementation, considering the energy functional of the 2-RDM and the case where $w_i = w$ for all i , we obtain the following bi-objective semidefinite program (SDP), which includes the relaxation for rank *via* the nuclear norm,

$$\begin{aligned} & \min_{{}^2D} E[{}^2D] + w \text{Tr}(\mathbf{E}_1 + \mathbf{E}_2) \\ \text{such that } & {}^2D \geq 0 \\ & {}^2Q \geq 0 \\ & {}^2G \geq 0 \\ & \mathbf{E}_1 \geq 0 \\ & \mathbf{E}_2 \geq 0 \\ & w > 0 \\ & \text{Tr}({}^2D) = N(N-1) \\ & {}^2Q = f_Q({}^2D) \\ & {}^2G = f_G({}^2D) \\ & {}^2\tilde{D} = {}^2D + \mathbf{E}_1 - \mathbf{E}_2 \\ & \tilde{S}_{pq}^n = \left((U \otimes U)^2 \tilde{D} (U \otimes U)^T \right)_{pq}^{pq}. \end{aligned} \quad (11)$$



The two error matrices, E_1 and E_2 , act as a “positive” and “negative” direction of error, respectively, increasing and decreasing the correction to the 2-RDM simultaneously. In the limit where $w \rightarrow 0$, the SDP reduces to the variational 2-RDM algorithm,^{33,39} in which only the energy functional $E[{}^2D]$ is minimized subject to the N -representability conditions. In the limit $w \rightarrow \infty$, the SDP gains the ability to treat arbitrary quantum states, including excited and non-stationary electronic states, within the same variational framework.^{77–79} In this work, we adopt equal weighting between the energy and measurement-consistency terms as a simple, minimally-biased default that avoids over-committing to either objective in the presence of noise. More generally, the w parameter may be guided by standard trade-off criteria (e.g., balancing energetic improvement against deviation from the measured shadow data). Additionally, a closely related study by the authors⁷⁷ shows that reconstruction performance is often robust across a fairly broad intermediate range of w , with clear improvement over a wide range of choices.

The present algorithm enforces a necessary subset of the N -representability conditions, known as the 2-positivity (DQG) conditions.^{41,56,57,59,80,81} The 2-positivity conditions restrict the particle–particle RDM (2D), the hole–hole RDM (2Q), and the particle–hole RDM (2G) to be positive semidefinite where f_Q and f_G are linear maps connecting 2Q and 2G to 2D , respectively. These semidefinite constraints correspond to keeping the particle–particle, hole–hole, and particle–hole probability distributions nonnegative. The 2-positivity conditions incur a computational scaling of $\mathcal{O}(r^4)$ in memory and $\mathcal{O}(r^6)$ in floating-point operations,³⁹ where r denotes the number of orbitals, and therefore, constrained shadow tomography maintains overall polynomial scaling. This scaling is consistent with variational 2-RDM SDP theory, where large-scale calculations have previously been demonstrated with millions of optimization variables (explicit timing data on standard hardware can be found in ref. 39). Scalability can further be modified by exploiting symmetry and solver structure, and partial constraint hierarchies (e.g., D or D_Q). Performing full 2-RDM tomography on a quantum device would nominally require $\mathcal{O}(r^4)$ distinct measurements, whereas our shadow-based technique reduces this requirement to approximately $n_s r^2$, where n_s is the number of classical shadows and typically satisfies $n_s \ll r^2$ for a desired accuracy ε .^{28,29}

From a practical standpoint, the error-mitigating formulation presented is particularly well-suited for application on quantum devices, where finite sampling noise and gate errors can corrupt 2-RDM estimates derived from shadow tomography or other measurement schemes. By incorporating error matrices directly into variational optimization, the method provides a mechanism to regularize noisy experimental data while maintaining physical consistency through the N -representability constraints. Beyond molecular electronic structure, this framework is broadly applicable to quantum simulation tasks involving spin models, correlated materials, and quantum state tomography, where convex relaxation and data-driven corrections offer a scalable and physically grounded route to extracting reliable observables from imperfect measurements.

3. Results

To implement constrained shadow tomography, we solve the semidefinite program (SDP) in eqn (11) by adding shadow constraints and nuclear norm relaxation to the variational 2-RDM (v2RDM) method^{33,82} in the Maple Quantum Chemistry Package.^{83,84} We refer to the resulting algorithm as the shadow v2RDM (sv2RDM) method. The SDP is solved using the boundary-point algorithm in ref. 39. To simulate measurement uncertainty, we generate noisy 2-RDMs by adding element-wise Gaussian perturbations to the full configuration interaction (FCI) 2-RDMs, corresponding to the Bernoulli shot noise expected from finite sampling (details can be found in the SI Sec. I A). The Fermionic Classical Shadows (FCS) protocol is implemented as described in ref. 24 (details can be found in the SI Sec. I B). Constrained shadow tomography is applied to the dissociation of the nitrogen dimer, ground-state energies of strongly correlated hydrogen chains⁸⁵ with up to ten equally spaced atoms, and a rectangle-square-rectangle transition of H_4 . The hydrogen atoms are represented in the minimal Slater-type orbital (STO-3G) basis set⁸⁶ while nitrogen is represented in the correlation-consistent polarized valence double-zeta (cc-pVDZ) basis set⁸⁷ in a 10 electrons-in-8 orbitals [10,8] active space. All quantum circuits used for data collection are implemented using the Qiskit package⁸⁸ and ffsim library⁸⁹ on IBM’s 156-qubit superconducting processor `ibm_fez` (details can be found in SI Sec. II).⁹⁰

We begin by assessing the accuracy and consistency of the sv2RDM approach against the most-favorably scaling FCS tomography protocol.²⁴ Fig. 1 summarizes the dependence of absolute energy error, computed in the minimal basis set, and the Frobenius norm of the 2-RDM error (relative to the 2-RDM from full configuration interaction (FCI)) on system size under fixed total shot budgets for hydrogen chains of increasing length (H_4 – H_{10}). Two sampling strategies were compared: the many-unitary/single-shot optimum for the FCS and the few-unitary/multiple-shot protocol for the sv2RDM. Each data point represents an average over 20 independent runs (10 for H_{10}), with 95% confidence intervals indicated by the error bars. Both quantum tomography methods are compared with the classical v2RDM method. Across all systems, the overall trends reveal sv2RDM outperforming FCS in both total energy and 2-RDM reconstruction with respect to the FCI results, given the same total shot budget for the two protocols. Importantly, the circuits corresponding to the shot budgets used are up to two orders of magnitude lower in depth for sv2RDM, substantially improving the feasibility of constrained shadow tomography on quantum hardware. Moreover, the sv2RDM method provides greater accuracy than the classical v2RDM method with increasing system size.

Fig. 2 shows the lowest eigenvalue of the 2-RDM as a function of the total shot budget for H_4 in the minimal basis, comparing the sv2RDM method with full 2-positivity conditions (DQG) to the FCS approach, which does not inherently produce N -representable matrices. As the total number of shots increases, both methods converge toward consistent eigenvalue



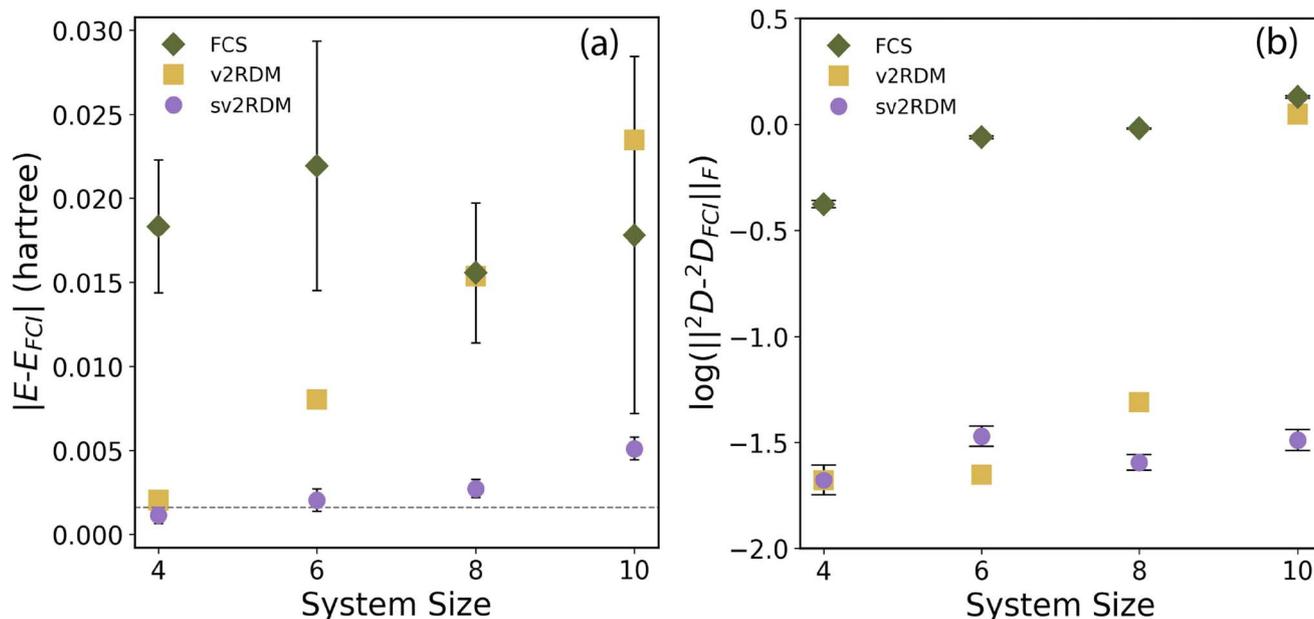


Fig. 1 (a) Absolute energy error as a function of system size under a fixed total shot budget. (b) Absolute Frobenius norm of the 2-RDM error (relative to 2-RDM from FCI) plotted on a logarithmic scale. Shot budgets compared are: 16 000 unitaries \times 1 shot vs. 16 unitaries \times 1000 shots for H_4 ; 36 000 unitaries \times 1 shot vs. 36 unitaries \times 1000 shots for H_6 ; 160 000 unitaries \times 1 shot vs. 160 unitaries \times 1000 shots for H_8 ; and 300 000 unitaries \times 1 shot vs. 300 unitaries \times 1000 shots for H_{10} . Each data point for the Fermionic Classical Shadows (FCS) and shadow variational 2-RDM (sv2RDM) methods represents the mean of 20 independent runs (10 for H_{10}). Error bars denote 95% confidence intervals (approximately 2σ) computed from these measurements. All calculations are performed using the minimal basis set.

estimates, indicating improved reconstruction of the underlying fermionic state. At low shot budgets, the FCS results exhibit larger fluctuations and unphysical negative eigenvalues, reflecting the stochastic nature of randomized measurement sampling. As expected from the scaling analysis in ref. 24, the statistical fluctuations in the FCS protocol diminish

approximately as $1/\sqrt{N_{\text{shots}}}$, leading to a smooth convergence toward the physically valid spectrum at greater sampling depths. In contrast, the sv2RDM (DQG) results remain stable across all sampling regimes with the most “negative” eigenvalue being zero, demonstrating that explicit enforcement of N -representability constraints significantly mitigates noise sensitivity and maintains physically valid 2-RDM spectra.

To further address the physicality of the reconstructed RDMs, Fig. 3 compares the absolute error matrices of the 2-RDMs obtained from sv2RDM (DQG) (a) and FCS (b) for the H_4 system in the minimal basis. Each matrix element represents the absolute deviation from the FCI reference value. The sv2RDM (DQG) reconstruction exhibits uniformly low and smoothly distributed errors, consistent with the enforcement of N -representability conditions. In contrast, the FCS results display more pronounced and localized deviations, across all elements, further highlighting the statistical noise associated with such stochastic measurement sampling. The visual contrast between the two panels emphasizes that the constrained sv2RDM approach yields systematically more accurate and physically consistent 2-RDMs than the unconstrained FCS method under comparable measurement resources.

Table 1 summarizes the Frobenius-norm differences between reconstructed and FCI ref. 2-RDMs across hydrogen chains of increasing length using the same total shot budgets. The SDP used in this case has $w \rightarrow \infty$ in eqn (11). The results, including only the D condition and the full DQG constraint set, are compared against the FCS protocol. Across all system sizes, imposing N -representability conditions systematically reduces

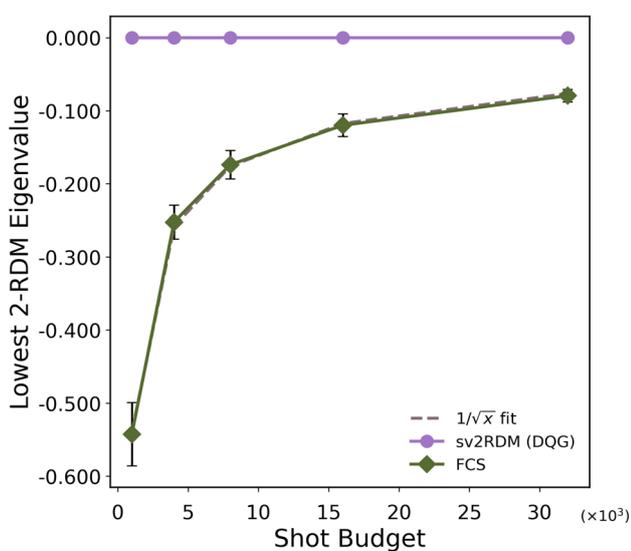


Fig. 2 Lowest 2-RDM eigenvalue as a function of total shot budget for the shadow variational 2-RDM (sv2RDM) method with full 2-positivity (DQG) constraints and for the Fermionic Classical Shadows (FCS) approach applied to H_4 in the minimal basis set.



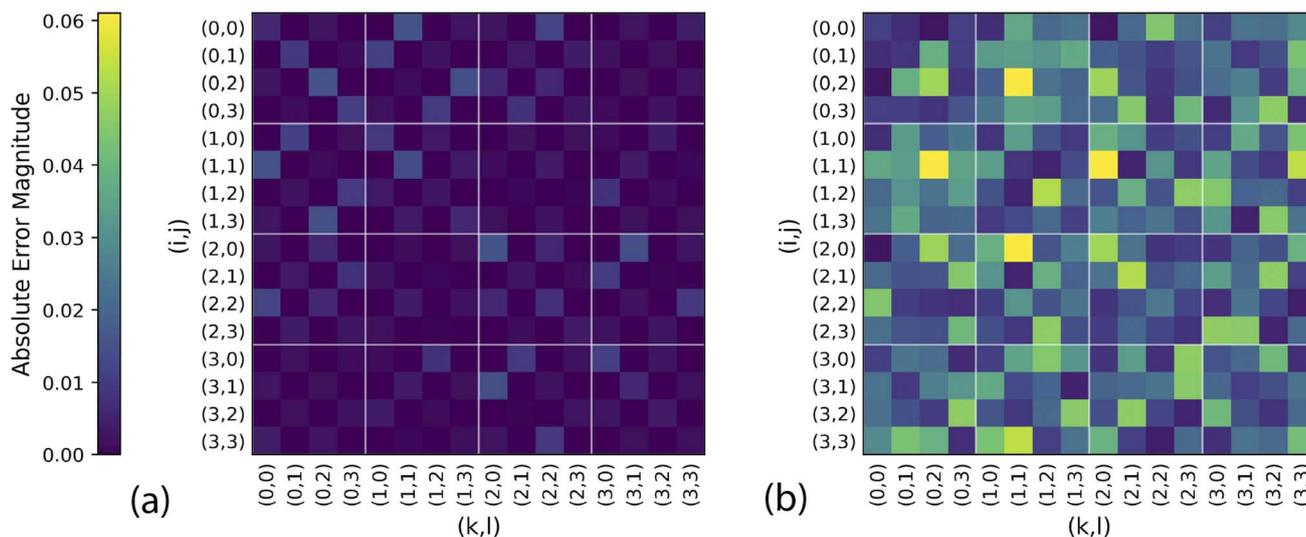


Fig. 3 Comparison of absolute error matrices for the H_4 system in the minimal basis set obtained using the shadow variational 2-RDM (sv2RDM) (a) and Fermionic Classical Shadows (FCS) (b) methods. Each matrix element represents the absolute deviation of the reconstructed 2-RDM from the FCI reference. The color intensity denotes the error magnitude, with brighter regions corresponding to larger deviations from the reference values.

Table 1 Comparison of Frobenius-norm differences with respect to the FCI 2-RDMs for various H chain lengths with the given total shot budgets when the semidefinite 2-RDM is optimized in the limit where $w \rightarrow \infty$ in eqn (11), with partial (D) and all 2-positivity conditions (DQG) included. Shot budgets compared are: 16 000 unitaries/1 shot vs. 16 unitaries/1000 shots for H_4 , 36 000 unitaries/1 shot vs. 36 unitaries/1000 shots for H_6 , 160 000 unitaries/1 shot vs. 160 unitaries/1000 shots for H_8 , and 300 000 unitaries/1 shot vs. 300 unitaries/1000 shots for H_{10}

H chain	Shot budget	D	DQG	FCS
4	16 000	0.129	0.059	0.423
6	36 000	0.226	0.080	0.873
8	160 000	0.454	0.037	0.960
10	300 000	0.217	0.055	1.349

the deviation from the reference state, with the greatest improvements observed when all DQG constraints are enforced. The effect becomes more pronounced for larger systems, where unconstrained or partially constrained optimizations show growing discrepancies from the FCS method. These results demonstrate that the inclusion of full 2-positivity conditions provides superior reconstruction fidelity, particularly under limited measurement resources and increasing system complexity.

The inclusion of electronic energy minimization acts as a Hamiltonian-informed regularizer, which is particularly beneficial in molecular ground-state settings, but is not required for the reconstruction itself. In the limit $w \rightarrow \infty$, the SDP reconstruction becomes fully independent of the Hamiltonian and may be applied to arbitrary quantum states. In this regime, we find that expectation values such as the electronic energy converge more slowly than norm-based 2-RDM metrics, consistent with the heightened sensitivity of energy

observables. Representative equilibrium-point energy benchmarks are provided in the SI (Table S1). Related work by the authors⁷⁷ further discusses the $w \rightarrow \infty$ limit as a useful setting for correlated purification and error mitigation, including applications beyond ground states such as excited-state reconstruction.

Fig. 4 presents the simulated potential energy curve of N_2 in the cc-pVDZ basis set with a [10,8] active space using complete active space configuration interaction (CASCI), sv2RDM, and FCS methods (SI Sec. III includes v2RDM results). The CASCI results serve as the reference for evaluating the accuracy of the reconstructed energy profiles. The FCS results exhibit noticeable deviations near the equilibrium and bond-stretching region, while the sv2RDM is able to closely reproduce reference results across the full dissociation range. These comparisons highlight that constrained semidefinite reconstructions provide quantitatively reliable potential energy surfaces even under limited measurement resources, whereas unconstrained fermionic shadow tomography approaches can struggle to maintain accuracy in strongly correlated regimes.

To evaluate the performance of constrained shadow tomography on real quantum hardware, the H_4 rectangle-square-rectangle potential energy profile is computed in the minimal basis using both classical and hybrid quantum approaches (Fig. 5). Density Functional Theory (DFT) results using the B3LYP functional^{91,92} are shown as the classical baseline, while the sv2RDM calculations are performed using measurements collected on the *ibm_fez* quantum processor (details can be found in SI Section II A). The Local Unitary Cluster Jastrow (LUCJ) ansatz⁹³ is optimized and computed classically, serving as the initially prepared quantum state and a reference for the sv2RDM calculations. The DFT results fail to reproduce the smooth energy flattening near the square geometry that signals



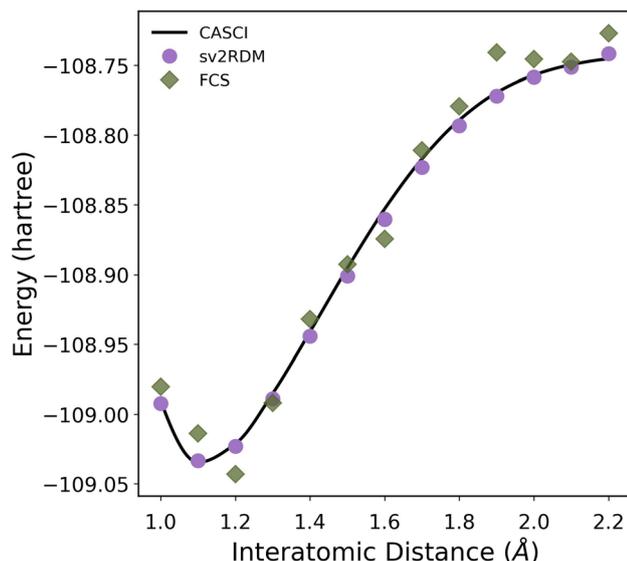


Fig. 4 Potential energy curve of N_2 computed using the cc-pVDZ basis set and a [10,8] active space with complete active space configuration interaction (CASCI), shadow variational 2-RDM (sv2RDM), and Fermionic Classical Shadows (FCS) methods. For sv2RDM and FCS, the total shot budgets used are 100 unitaries \times 1000 shots and 100 000 unitaries \times 1 shot, respectively.

strong static correlation. In contrast, the sv2RDM results recover the correct qualitative features, closely matching the expected multireference character. The sv2RDM energies obtained from the quantum device align well with the LUCJ benchmark within statistical uncertainty, demonstrating that

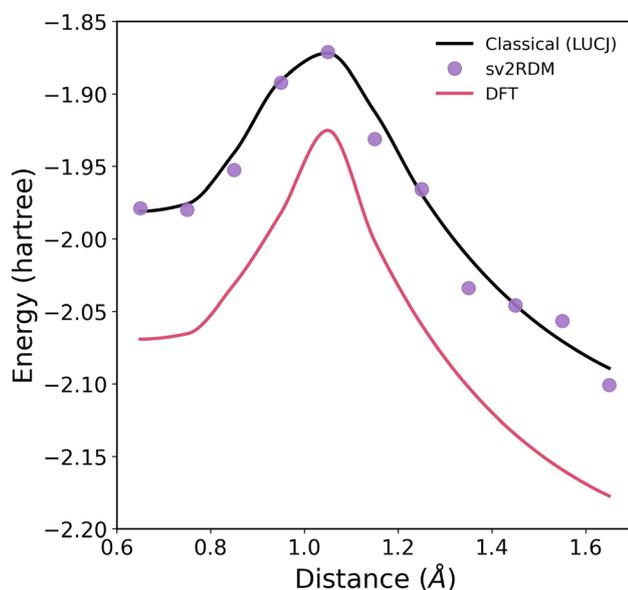


Fig. 5 Potential energy curve for the H_4 rectangle–square–rectangle transition in the minimal basis set, computed using Density Functional Theory (DFT) with the B3LYP functional as a classical baseline, the classically computed Local Unitary Cluster Jastrow (LUCJ) ansatz, and the shadow variational 2-RDM (sv2RDM) method computed using measurements collected on the ibm_fez quantum computer.

Table 2 Comparison of correlation energies with respect to Hartree Fock for various H chain lengths. Measurements for sv2RDM calculations are collected on the ibm_fez quantum device with the following shot budgets: 16 unitaries/10 000 shots for H_4 , 36 unitaries/10 000 shots for H_6 , 160 unitaries/10 000 shots for H_8 , and 300 unitaries/10 000 shots for H_{10} . Hartree Fock, LUCJ, and FCI calculations are run classically. All calculations use a minimal basis set

H chain	sv2RDM	LUCJ	FCI
4	0.011	0.034	0.068
6	0.013	0.062	0.101
8	0.024	0.061	0.133
10	0.065	0.084	0.166

error-mitigated semidefinite constraints can effectively suppress noise-induced artifacts and maintain physical consistency in near-term quantum simulations.

The scalability of the constrained reconstruction approach is further evaluated by comparing correlation energies for hydrogen chains of increasing length, computed with sv2RDM, FCI, and the LUCJ ansatz in the minimal basis (Table 2). sv2RDM energies are obtained from measurement data collected on the ibm_fez quantum processor under fixed total shot budgets scaled with system size, with up to 20 qubits (details can be found in SI Section II B). Other methods were computed classically. Across all chain lengths, sv2RDM reproduces the qualitative energy trends of the reference LUCJ results, maintaining size-extensive behavior and smooth energy variation. Though deviations from the classical references appear at all system sizes, attributable to the increased noise inherent to hardware sampling, the results highlight that, even with limited quantum resources, enforcing semidefinite constraints enables systematically scalable energy predictions for chemical systems.

The results demonstrate that constrained shadow tomography provides a robust and systematically improvable framework for extracting physically meaningful quantities from both simulated and experimental quantum data. The numerical benchmarks employ Gaussian perturbations as a controlled model of finite-shot statistical noise, enabling a clear assessment of the proposed tomography's ability to suppress statistical fluctuations and reduce the number of required measurements. While this noise model does not capture realistic experimental effects, including coherent gate errors, readout bias, and correlated device effects, these are naturally present in the accompanying hardware demonstrations, which show qualitatively similar behavior. This agreement supports the robustness and measurement efficiency of the method beyond idealized shot-noise assumptions.

Across molecular benchmarks, enforcing N -representability through the 2-positivity conditions yields stable energy estimates, well-behaved eigenvalue spectra, and accurate two-particle correlations even under modest shot budgets. Comparisons with classical v2RDM, LUCJ, and FCI references confirm that the sv2RDM approach retains quantitative accuracy while remaining resilient to hardware noise and stochastic sampling errors. The constrained shadow tomography is



complementary to established quantum error mitigation techniques such as symmetry verification⁹⁴ and purification-based post-processing,⁹⁵ which enforce physical structure at the level of the global state or density matrix. By contrast, the present approach operates directly at the level of fermionic reduced density matrices, enforcing N -representability constraints through a convex SDP, and thus provides a noise-model-independent route to suppressing unphysical errors in measured observables. This framework is naturally compatible with other mitigation layers, including correlated purification strategies,⁷⁷ as well as probabilistic error cancellation methods⁹⁶ that rely on explicit noise characterization. These findings establish a consistent foundation for assessing scalability, error mitigation, and resource efficiency in near-term quantum simulations, motivating a deeper discussion on methodological implications and potential pathways toward chemically accurate hybrid quantum-classical frameworks.

4. Discussion and conclusions

The results presented in this work highlight the advantages of incorporating physical constraints into scalable quantum state reconstruction protocols. By embedding N -representability directly into the shadow tomography framework, our constrained shadow tomography formulation enables physically consistent and noise-resilient reconstructions of RDMs. The integration of bi-objective semidefinite programming, nuclear-norm regularization, and N -representability jointly suppresses measurement-induced errors, yielding accurate and stable estimates under limited shot budgets and realistic noise conditions. Benchmark results for hydrogen chains and the N_2 molecule confirm that enforcing physical constraints during reconstruction improves both accuracy and robustness across a range of system sizes and sampling budgets, while remaining compatible with shallow circuits. Beyond quantitative improvements, this approach provides a unified and computationally efficient framework that connects statistical inference, convex optimization, and quantum chemistry post-processing. The demonstrated reduction in measurement requirements and enhanced robustness underscore the potential of constraint-enforced tomography as a key component of hybrid quantum-classical workflows for benchmarking and validating quantum simulations of correlated fermionic systems.

The presented algorithm establishes a practical foundation for advancing quantum state characterization in many-body quantum simulations. Future work may explore the incorporation of additional system symmetries or adaptive weighting strategies to further enhance RDM reconstruction fidelity. As quantum hardware continues to progress, methods that integrate physical insight, optimization principles, and statistical learning will be essential for developing reliable and scalable approaches to quantum state reconstruction. This study positions constrained shadow tomography as a powerful and generalizable framework for extracting accurate physical insight from limited quantum data, paving the way for robust, physically grounded characterization of complex quantum systems.

Author contributions

Conceptualization and methodology: I. A. and D. A. M.; investigation, data curation, formal analysis, software, visualization: all authors; supervision, funding acquisition, project administration, resources, validation: K. J. S. and D. A. M.; writing – original draft: I. A.; writing – review & editing: all authors.

Conflicts of interest

There are no conflicts to declare.

Data availability

All data and sample computational script are openly available at: <https://github.com/damaz/ConstrainedShadowTomography.git>. The data supporting this article have been included as part of the supplementary information (SI). Supplementary information is available. See DOI: <https://doi.org/10.1039/d5sc09162d>.

Acknowledgements

We thank Nate Earnest-Noble and Jamie Garcia for valuable discussions. D. A. M. gratefully acknowledges support from IBM under the IBM-UChicago Quantum Collaboration. D. A. M. also gratefully acknowledges support from the U.S. National Science Foundation under Grant No. CHE-2155082. I. A. gratefully acknowledges the NSF Graduate Research Fellowship Program under Grant No. 2140001.

References

- 1 K. Vogel and H. Risken, Determination of quasiprobability distributions in terms of probability distributions for the rotated quadrature phase, *Phys. Rev. A: At., Mol., Opt. Phys.*, 1989, **40**, 2847–2849.
- 2 Z. Hradil, Quantum-state estimation, *Phys. Rev. A: At., Mol., Opt. Phys.*, 1997, **55**, R1561–R1564.
- 3 D. F. V. James, P. G. Kwiat, W. J. Munro and A. G. White, Measurement of qubits, *Phys. Rev. A: At., Mol., Opt. Phys.*, 2001, **64**, 052312.
- 4 *Quantum State Estimation, Lecture Notes in Physics*, ed. M. Paris and J. Řeháček, Springer, Berlin, Heidelberg, 2004, vol. 649.
- 5 K. Head-Marsden, J. Flick, C. J. Ciccarino and P. Narang, Quantum Information and Algorithms for Correlated Quantum Matter, *Chem. Rev.*, 2021, **121**, 5.
- 6 C. L. Benavides-Riveros, L. Chen, C. Schilling, S. Mantilla and S. Pittalis, Excitations of Quantum Many-Body Systems via Purified Ensembles: A Unitary-Coupled-Cluster-Based Approach, *Phys. Rev. Lett.*, 2022, **129**, 2201–10974.
- 7 S. McArdle, S. Endo, A. Aspuru-Guzik, S. C. Benjamin and X. Yuan, Quantum computational chemistry, *Rev. Mod. Phys.*, 2020, **92**, 015003.
- 8 S. E. Smart and D. A. Mazziotti, Quantum solver of contracted eigenvalue equations for scalable molecular



- simulations on quantum computing devices, *Phys. Rev. Lett.*, 2021, **126**, 070504.
- 9 D. L. Donoho, Compressed sensing, *IEEE Trans. Inf. Theor.*, 2006, **52**, 1289.
- 10 E. J. Candès, J. K. Romberg and T. Tao, Stable signal recovery from incomplete and inaccurate measurements, *Commun. Pure Appl. Math.*, 2006, **59**, 1207–1223.
- 11 J. Carrasquilla, G. Torlai, R. G. Melko and L. Aolita, Reconstructing quantum states with generative models, *Nat. Mach. Intell.*, 2019, **1**, 155–161.
- 12 K. Mukherjee, J. Schachenmayer, S. Whitlock and S. Wüster, Quantum network tomography of small rydberg arrays by machine learning, *Phys. Rev. A*, 2025, **111**, 043319.
- 13 S. T. Flammia and Y.-K. Liu, Direct fidelity estimation from few pauli measurements, *Phys. Rev. Lett.*, 2011, **106**, 230501.
- 14 X. Zhang, M. Luo, Z. Wen, Q. Feng, S. Pang, W. Luo and X. Zhou, Direct fidelity estimation of quantum states using machine learning, *Phys. Rev. Lett.*, 2021, **127**, 130503.
- 15 A. Seshadri, M. Ringbauer, J. Spainhour, R. Blatt, T. Monz and S. Becker, Versatile fidelity estimation with confidence, *Phys. Rev. Lett.*, 2024, **133**, 020402.
- 16 V. Singh, R. Gupta, M. Sajjan, F. Remacle, R. D. Levine and S. Kais, Maximal entropy formalism and the restricted boltzmann machine, *J. Phys. Chem. A*, 2025, **129**, 5405–5414.
- 17 S. Aaronson, Shadow tomography of quantum states, *SIAM J. Comput.*, 2020, **49**, STOC18.
- 18 H.-Y. Huang, R. Kueng and J. Preskill, Predicting many properties of a quantum system from very few measurements, *Nat. Phys.*, 2020, **16**, 1050.
- 19 D.-J. Zhang and D. M. Tong, Krylov shadow tomography: Efficient estimation of quantum fisher information, *Phys. Rev. Lett.*, 2025, **134**, 110802.
- 20 J. Castaneda and N. Wiebe, Hamiltonian learning via shadow tomography of pseudo-choi states, *Quantum*, 2025, **9**, 1700.
- 21 R. D. Somma, R. King, R. Kothari, T. E. O'Brien and R. Babbush, Shadow hamiltonian simulation, *Nat. Commun.*, 2025, **16**, 57451.
- 22 R. King, D. Gosset, R. Kothari, and R. Babbush, Triply efficient shadow tomography, in *Proceedings of the 2025 Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*, Society for Industrial and Applied Mathematics, 2025, pp. 914–946.
- 23 A. Zhao, N. C. Rubin and A. Miyake, Fermionic partial tomography via classical shadows, *Phys. Rev. Lett.*, 2021, **127**, 110504.
- 24 G. H. Low, Classical Shadows of Fermions with Particle Number Symmetry, *arXiv*, 2022, preprint, arXiv:2208.08964, DOI: [10.48550/arXiv.2208.08964](https://doi.org/10.48550/arXiv.2208.08964).
- 25 W. J. Huggins, B. A. O'Gorman, N. C. Rubin, D. R. Reichman, R. Babbush and J. Lee, Unbiasing fermionic quantum monte carlo with a quantum computer, *Nature*, 2022, **603**, 416–420.
- 26 A. Zhao and A. Miyake, Group-theoretic error mitigation enabled by classical shadows and symmetries, *Npj Quantum Inf.*, 2024, **10**, 854.
- 27 B. Huang, Y.-T. Chen, B. Gupt, M. Suchara, A. Tran, S. McArdle and G. Galli, Evaluating a quantum-classical quantum monte carlo algorithm with matchgate shadows, *Phys. Rev. Res.*, 2024, **6**, 043063.
- 28 I. Avdic and D. A. Mazziotti, Fewer measurements from shadow tomography with N -representability conditions, *Phys. Rev. Lett.*, 2024, **132**, 220802.
- 29 I. Avdic and D. A. Mazziotti, Enhanced shadow tomography of molecular excited states via the enforcement of N -representability conditions by semidefinite programming, *Phys. Rev. A*, 2024, **110**, 052407.
- 30 D. A. Mazziotti and R. M. Erdahl, Uncertainty relations and reduced density matrices: Mapping many-body quantum mechanics onto four particles, *Phys. Rev. A: At., Mol., Opt. Phys.*, 2001, **63**, 042113.
- 31 M. Nakata, H. Nakatsuji, M. Ehara, M. Fukuda, K. Nakata and K. Fujisawa, Variational calculations of fermion second-order reduced density matrices by semidefinite programming algorithm, *J. Chem. Phys.*, 2001, **114**, 8282.
- 32 D. A. Mazziotti, Variational minimization of atomic and molecular ground-state energies via the two-particle reduced density matrix, *Phys. Rev. A: At., Mol., Opt. Phys.*, 2002, **65**, 062511.
- 33 D. A. Mazziotti, Realization of Quantum Chemistry without Wave Functions through First-Order Semidefinite Programming, *Phys. Rev. Lett.*, 2004, **93**, 213001.
- 34 Z. Zhao, B. J. Braams, M. Fukuda, M. L. Overton and J. K. Percus, The reduced density matrix method for electronic structure calculations and the role of three-index representability conditions, *J. Chem. Phys.*, 2004, **120**, 2095.
- 35 E. Cancès, G. Stoltz and M. Lewin, The electronic ground-state energy problem: A new reduced density matrix approach, *J. Chem. Phys.*, 2006, **125**, 064101.
- 36 G. Gidofalvi and D. A. Mazziotti, Active-space two-electron reduced-density-matrix method: Complete active-space calculations without diagonalization of the N -electron Hamiltonian, *J. Chem. Phys.*, 2008, **129**, 134108.
- 37 N. Shenvi and A. F. Izmaylov, Active-Space N -Representability Constraints for Variational Two-Particle Reduced Density Matrix Calculations, *Phys. Rev. Lett.*, 2010, **105**, 213003.
- 38 B. Verstichel, H. v. Aggelen, W. Poelmans and D. V. Neck, Variational Two-Particle Density Matrix Calculation for the Hubbard Model Below Half Filling Using Spin-Adapted Lifting Conditions, *Phys. Rev. Lett.*, 2011, **108**, 213001.
- 39 D. A. Mazziotti, Large-scale semidefinite programming for many-electron quantum mechanics, *Phys. Rev. Lett.*, 2011, **106**, 083001.
- 40 T. Baumgratz and M. B. Plenio, Lower bounds for ground states of condensed matter systems, *New J. Phys.*, 2012, **14**, 11065275.
- 41 D. A. Mazziotti, Structure of fermionic density matrices: Complete N -representability conditions, *Phys. Rev. Lett.*, 2012, **108**, 263002.
- 42 D. A. Mazziotti, Enhanced Constraints for Accurate Lower Bounds on Many-Electron Quantum Energies from Variational Two-Electron Reduced Density Matrix Theory, *Phys. Rev. Lett.*, 2016, **117**, 153001.



- 43 D. R. Alcoba, A. Torre, L. Lain, G. E. Massaccesi, O. B. Oña, E. M. Honoré, W. Poelmans, D. V. Neck, P. Bultinck and S. D. Baerdemacker, Direct variational determination of the two-electron reduced density matrix for doubly occupied-configuration-interaction wave functions: The influence of three-index N -representability conditions, *J. Chem. Phys.*, 2018, **148**, 024105.
- 44 D. A. Mazziotti, Dual-cone variational calculation of the two-electron reduced density matrix, *Phys. Rev. A*, 2020, **102**, 2103–17155.
- 45 R. R. Li, M. D. Liebenthal and A. E. DePrince, Challenges for variational reduced-density-matrix theory with three-particle N -representability conditions, *J. Chem. Phys.*, 2021, **155**, 174110.
- 46 M. J. Knight, H. M. Quiney and A. M. Martin, Reduced density matrix approach to ultracold few-fermion systems in one dimension, *New J. Phys.*, 2022, **24**, 053004.
- 47 D. A. Mazziotti, Quantum many-body theory from a solution of the N -representability problem, *Phys. Rev. Lett.*, 2023, **130**, 153001.
- 48 Q. Gao, R. A. Lanzetta, P. Ledwith, J. Wang and E. Khalaf, Bootstrapping the Quantum Hall Problem, *Phys. Rev. X*, 2025, **15**, 031034.
- 49 A. O. Schouten, S. Ewing, and D. A. Mazziotti, Bootstrapping the Electronic Structure of Quantum Materials, *arXiv*, 2025, preprint, arXiv:2504.02861, DOI: [10.48550/arXiv.2504.02861](https://doi.org/10.48550/arXiv.2504.02861).
- 50 M. Piris, Global Natural Orbital Functional: Towards the Complete Description of the Electron Correlation, *Phys. Rev. Lett.*, 2021, **127**, 2112–2119.
- 51 J. Haah, A. W. Harrow, Z. Ji, X. Wu and N. Yu, Sample-optimal tomography of quantum states, *IEEE Trans. Inf. Theor.*, 2017, **63**, 5628.
- 52 A. Elben, J. Yu, G. Zhu, M. Hafezi, F. Pollmann, P. Zoller and B. Vermersch, Many-body topological invariants from randomized measurements in synthetic quantum matter, *Sci. Adv.*, 2020, **6**, eaz3666.
- 53 A. Elben, S. T. Flammia, H.-Y. Huang, R. Kueng, J. Preskill, B. Vermersch and P. Zoller, The randomized measurement toolbox, *Nat. Rev. Phys.*, 2022, **5**, 9.
- 54 S. B. Bravyi and A. Y. Kitaev, Fermionic quantum computation, *Ann. Phys.*, 2002, **298**, 210–226.
- 55 Z. Jiang, A. Kalev, W. Mruzckiewicz and H. Neven, Optimal fermion-to-qubit mapping via ternary trees with applications to reduced quantum states learning, *Quantum*, 2020, **4**, 276.
- 56 D. A. Mazziotti, Reduced-Density-Matrix Mechanics: With Application to Many-Electron Atoms and Molecule, *Adv. Chem. Phys.*, 2007, **134**, 3–11.
- 57 A. J. Coleman, Structure of Fermion Density Matrices, *Rev. Mod. Phys.*, 1963, **35**, 668.
- 58 A. Coleman and V. Yukalov, *Reduced Density Matrices: Coulson's Challenge*, *Lecture Notes in Chemistry*, Springer Berlin Heidelberg, 2000.
- 59 C. Garrod and J. K. Percus, Reduction of the N -Particle Variational Problem, *J. Math. Phys.*, 1964, **5**, 1756.
- 60 H. Kummer, N -representability problem for reduced density matrices, *J. Math. Phys.*, 1967, **8**, 2063.
- 61 R. M. Erdahl, Representability, *Int. J. Quantum Chem.*, 1978, **13**, 697.
- 62 D. A. Mazziotti, Variational reduced-density-matrix method using three-particle N -representability conditions with application to many-electron molecules, *Phys. Rev. A: At., Mol., Opt. Phys.*, 2006, **74**, 032501.
- 63 A. Haar, Der massbegriff in der theorie der kontinuierlichen gruppen, *Ann. Math.*, 1933, **34**, 147.
- 64 S. E. Smart, J.-N. Boyn and D. A. Mazziotti, Resolving correlated states of benzyne with an error-mitigated contracted quantum eigensolver, *Phys. Rev. A*, 2022, **105**, 022405.
- 65 E. J. Candès and B. Recht, Exact Matrix Completion via Convex Optimization, *Found. Math. Comput.*, 2009, **9**, 717.
- 66 J.-F. Cai, E. J. Cands and Z. Shen, A Singular Value Thresholding Algorithm for Matrix Completion, *SIAM J. Optim.*, 2010, **20**, 1956.
- 67 T. E. O'Brien, B. Senjean, R. Sagastizabal, X. Bonet-Monroig, A. Dutkiewicz, F. Buda, L. DiCarlo and L. Visscher, Calculating energy derivatives for quantum chemistry on a quantum computer, *Npj Quantum Inf.*, 2019, **5**, 113.
- 68 C. Overy, G. H. Booth, N. S. Blunt, J. J. Shepherd, D. Cleland and A. Alavi, Unbiased reduced density matrices and electronic properties from full configuration interaction quantum monte carlo, *J. Chem. Phys.*, 2014, **141**, 244117.
- 69 G. Gidofalvi and D. A. Mazziotti, Molecular properties from variational reduced-density-matrix theory with three-particle N -representability conditions, *J. Chem. Phys.*, 2007, **126**, 024105.
- 70 J. P. Evans and R. E. Steuer, A revised simplex method for linear multiple objective programs, *Math. Program.*, 1973, **5**, 54.
- 71 M. Zeleny, *Linear Multi-Objective Programming*, Springer-Verlag, New York, NY, 1974.
- 72 J. J. Foley and D. A. Mazziotti, Measurement-driven reconstruction of many-particle quantum processes by semidefinite programming with application to photosynthetic light harvesting, *Phys. Rev. A: At., Mol., Opt. Phys.*, 2012, **86**, 012512.
- 73 A. Kotil, E. Pelofske, S. Riedmüller, D. J. Egger, S. Eidenbenz, T. Koch and S. Woerner, Quantum approximate multi-objective optimization, *Nat. Comput. Sci.*, 2025, **5**, 873.
- 74 B. Recht, M. Fazel and P. A. Parrilo, Guaranteed minimum-rank solutions of linear matrix equations via nuclear norm minimization, *SIAM Rev.*, 2010, **52**, 471–501.
- 75 C. Beck and R. D'Andrea, Computational study and comparisons of LFT reducibility methods, in *Proceedings of the 1998 American Control Conference (ACC)*, IEEE, 1998, vol. 2, pp. 1013–1017.
- 76 M. Mesbahi and G. Papavasilopoulos, On the rank minimization problem over a positive semidefinite linear matrix inequality, *IEEE Trans. Autom. Control*, 1997, **42**, 239–243.
- 77 Y. Wang, I. Avdic, M. Rose, L. I. P. Torres, A. O. Schouten, K. J. Sung and D. A. Mazziotti, *Correlated purification for restoring n -representability in quantum simulation*, 2025.



- 78 J. Liebert, A. O. Schouten, I. Avdic, C. Schilling and D. A. Mazziotti, *Refining ensemble n -representability of one-body density matrices from partial information*, 2025.
- 79 J. Liebert, F. Castillo, J.-P. Labbé and C. Schilling, Foundation of one-particle reduced density matrix functional theory for excited states, *J. Theor. Comput. Chem.*, 2021, **18**, 124.
- 80 D. A. Mazziotti and R. M. Erdahl, Uncertainty relations and reduced density matrices: Mapping many-body quantum mechanics onto four particles, *Phys. Rev. A: At., Mol., Opt. Phys.*, 2001, **63**, 042113.
- 81 C. Schilling, Quasipinning and its relevance for N -fermion quantum states, *Phys. Rev. A: At., Mol., Opt. Phys.*, 2015, **91**, 022105.
- 82 G. Gidofalvi and D. A. Mazziotti, Spin and symmetry adaptation of the variational two-electron reduced-density-matrix method, *Phys. Rev. A: At., Mol., Opt. Phys.*, 2005, **72**, 052505.
- 83 *Maple (Maplesoft, Waterloo)*, 2025.
- 84 RDMChem, *Quantum Chemistry Package, Maplesoft, Waterloo*, 2025.
- 85 S. Suhai, Electron correlation in extended systems: Fourth-order many-body perturbation theory and density-functional methods applied to an infinite chain of hydrogen atoms, *Phys. Rev. B: Condens. Matter Mater. Phys.*, 1994, **50**, 14791.
- 86 W. J. Hehre, R. F. Stewart and J. A. Pople, Self-consistent molecular-orbital methods. i. use of gaussian expansions of slater-type atomic orbitals, *Chem. Phys.*, 1969, **51**, 2657–2664.
- 87 T. H. Dunning, Gaussian basis sets for use in correlated molecular calculations. i. the atoms boron through neon and hydrogen, *Chem. Phys.*, 1989, **90**, 1007–1023.
- 88 Qiskit contributors, *Qiskit: An open-source framework for quantum computing*, 2025.
- 89 The ffsim developers, ffsim: Faster simulations of fermionic quantum circuits, <https://github.com/qiskit-community/ffsim>.
- 90 All data and sample computational script are openlyavailable at <https://github.com/damaz/ConstrainedShadowTomography>.
- 91 P. J. Stephens, F. J. Devlin, C. F. Chabalowski and M. J. Frisch, Ab initio calculation of vibrational absorption and circular dichroism spectra using density functional force fields, *J. Phys. Chem.*, 1994, **98**, 11623–11627.
- 92 A. D. Becke, Density-functional thermochemistry. iii. the role of exact exchange, *J. Chem. Phys.*, 1993, **98**, 5648–5652.
- 93 M. Motta, K. J. Sung, K. B. Whaley, M. Head-Gordon and J. Shee, Bridging physical intuition and hardware efficiency for correlated electronic states: the local unitary cluster jastrow ansatz for electronic structure, *Chem. Sci.*, 2023, **14**, 11213.
- 94 R. Sagastizabal, X. Bonet-Monroig, M. Singh, *et al.*, Experimental error mitigation via symmetry verification in a variational quantum eigensolver, *Phys. Rev. A: At., Mol., Opt. Phys.*, 2019, **100**, 010302.
- 95 T. E. O'Brien, G. Anselmetti, F. Gkritis, *et al.*, Purification-based quantum error mitigation of pair-correlated electron simulations, *Nat. Phys.*, 2023, **19**, 1787.
- 96 K. Temme, S. Bravyi and J. M. Gambetta, Error mitigation for short-depth quantum circuits, *Phys. Rev. Lett.*, 2017, **119**, 180509.

