

Classical optimization of reduced density matrix estimations with classical shadows using N -representability conditions under shot noise considerations

Gian-Luca R. Anselmetti ^{*}, Matthias Degroote , Nikolaj Moll , Raffaele Santagati , and Michael Streif 
Quantum Lab, *Boehringer Ingelheim*, 55218 Ingelheim am Rhein, Germany



(Received 24 October 2025; accepted 14 July 2026; published 10 August 2026)

Classical shadow tomography has become a powerful tool in learning about quantum states prepared on a quantum computer. Recent works have used classical shadows to variationally enforce N -representability conditions on the two-particle reduced density matrix. In this paper, we build upon previous research by choosing an improved estimator within classical shadow tomography and rephrasing the optimization constraints, resulting in an overall enhancement in performance under comparable measurement shot budgets. We further explore the specific regimes where these methods outperform the unbiased estimator of the stand-alone classical shadow protocol and quantify the potential savings in numerical studies.

DOI: [10.1103/hwcl-yfy4](https://doi.org/10.1103/hwcl-yfy4)

I. INTRODUCTION

The simulation of quantum chemistry presents a hard challenge for the field of computation. In recent years, quantum computation has been proposed to aid quantum chemistry, offering to simulate complex systems that previously were out of reach for conventional methods [1–4]. However, current quantum computers are still limited in size and the number of operations due to noise and the absence of error correction [5]. As such, stand-alone quantum algorithms competing with the highly optimized methods of traditional quantum chemistry are still out of reach for the foreseeable future. This has opened up an avenue of research at the intersection of purely classical methods with the support of newly proposed simulation techniques on a quantum computer [6,7].

Early efforts have sought to combine the strengths of classical and quantum algorithms in an approach labeled as hybrid quantum computation. In this spirit, recent work has tried to enhance some established classical methods with data gathered from shadow tomography, a powerful tool preparing efficient classical representations of quantum states prepared on a quantum computer for certain observables [8–11]. An example is the combination of this tool with traditional methods of quantum chemistry like quantum Monte Carlo [6] and coupled cluster [7], achieving improvement in accuracy in certain situations. These techniques are motivated by the fact that learning on quantum data in the form of classical shadows is known to provide quantum advantage [12] under certain conditions, identifying sampling on a quantum computer as a minimal-cost useful task.

While the aforementioned methods have focused on enhancing classical methods with quantum data, other work has sought to apply techniques developed in classical quantum chemistry to quantum computation, e.g., constraints stemming from the N -representability conditions [limiting the space of valid reduced density matrices (RDM) that can stem from a physical wave function] to enhance quantum readout [13–15]. The goal of these works is to optimize measurement methods in sampling from an unknown quantum state, mostly trying to decrease sample variance and reduce sensitivity to device error [13,14] and to decrease the shot budget of shadow tomography [15] with the use of v2rdm (variational 2-RDM) [16,17], a method casting the N -representability conditions into a semidefinite program. To achieve industrial relevance, extracting ground-state energies from quantum calculation does not lead to an application *per se*. The promise lies more in the efficient extraction of further chemically relevant observables, like, e.g., molecular forces for geometry optimizations or molecular dynamics [18,19], as this is paramount to compete with established classical methods [20]. As fault-tolerant algorithms require a large overhead to extract expectation values of observables on ground states [21,22], sampling shadows of ground states from a quantum computer could help extract reduced density matrices more efficiently from a v2rdm procedure.

In this article, we introduce a way of constraining the semidefinite program by using an improved classical shadow estimator of the entire 2-RDM estimate. We also further explore the role of shot noise in estimation problems with and without the use of v2rdm in the classical shadow tomography. We establish regimes where the method can lead to savings in shot budgets of up to a factor of 15 over the nonoptimized estimate. We do not find that previous approaches lead to an improvement when accounting for shot noise and further explore behavior when choosing different cost functions.

The article is structured in the following way: First, we give a brief introduction to shadow tomography (Sec. II) and v2rdm (Sec. III). The former is focused on fermionic

^{*}Contact author: gian-luca.anselmetti@boehringer-ingelheim.com

tomography respecting particle number and spin preservation, and the latter briefly introduces core concepts of the constraints respected in our optimization and how to deal with inequalities when casting semidefinite programs. We then describe the two additional constraints derived from quantum data (Sec. IV), which we consider in our numerics, one from previous work [15] and one introduced in this work. We then simulate performance for various systems (Sec. V) and compare it with the pure classical estimate and the pure quantum estimate derived from shadow tomography without reoptimizing the 2-RDM to respect the N -representability conditions.

II. CLASSICAL SHADOW PROTOCOL

Classical shadows, a scheme introduced in Ref. [8] based on the shadow tomography proposal [23], can predict many properties of a quantum system with a number of measurements that scale logarithmically with the number of properties measured. The task of the protocol is to estimate the expectation value of multiple observables $\langle O_i \rangle$ on an unknown quantum state ρ . This requires the preparation of multiple copies of the state measured under random unitaries U drawn from an ensemble $\mathcal{U}$, resulting in the measurement outcome $|b\rangle$. Averaging over the results constitutes the measurement channel

$$\mathcal{M}(\rho) = \mathbb{E}_{U \sim \mathcal{U}} \left[\sum_{b \in \{0,1\}^N} \langle b|U\rho U^\dagger|b\rangle U^\dagger|b\rangle\langle b|U \right]. \quad (1)$$

The collection of random unitaries U and associated measurement outcomes $\langle b|$ is referred to as a classical shadow. For each of these pairs, the measurement channel can be inverted to build an unbiased estimator of the quantum state

$$\hat{\rho} = \mathcal{M}^{-1}(U^\dagger|b\rangle\langle b|U). \quad (2)$$

Throughout this work, hatted quantities $\hat{\rho}$ denote statistical estimators of the corresponding true quantities ρ . From this classical estimator, one can infer the expected values of the desired observables $\langle O_i \rangle = \text{Tr}(O_i \rho)$. The quality of the inverse of the measurement channel $\mathcal{M}^{-1}$ and the efficiency of the classical and quantum procedures depend on the chosen ensemble of unitaries and the associated observables one estimates from the quantum state.

Multiple choices for the ensembles and associated classical estimation schemes have been proposed [9–11]. We focus here on shadow estimation schemes twirling over the particle number restricted ensemble of matchgates [11]. This constitutes an efficient estimation scheme independent of the number of qubits, N , and its quantum cost scales only in the number of occupied fermionic modes η , which is the particle number of the trial state. This ensemble is usually referred to as the ensemble of single-particle basis rotations or orbital rotations $U(u)$, which define all valid transformations between different sets of the fermionic creation and annihilation operators a_q and $a_q^\dagger$ by

$$U^\dagger(u)a_p^\dagger U(u) = \sum_{q=1}^N u_{pq}a_q^\dagger, \quad (3)$$

where $u \in \mathbb{C}^{N \times N}$ is a unitary matrix defining the change of basis of the fermionic modes, with their representation $U(u) \in \mathbb{C}^{2^N \times 2^N}$ in Hilbert space. We refer to the submatrix restricted to the k -particle subspace as ${}^kU \in \mathbb{C}^{\binom{N}{k} \times \binom{N}{k}}$, and therefore ${}^2U$ as the restriction to the two-particle subspace. We further restrict unitaries beyond preserving particle number to also preserve spin in Appendix C as these are the most important symmetries for molecular systems.

In general, expectation values of generic k -particle reduced density matrices can be inferred from this procedure [11], we will only handle the two ($k=2$)-particle reduced density matrices as no higher-order density matrices are required for the purpose of estimating the observables of our interest. Our task, therefore, is to estimate the quantity

$$\langle {}_S^2 \hat{\mathbf{D}}^{pqrs} \rangle = \text{Tr}[a_p^\dagger a_q^\dagger a_s a_r \rho], \quad (4)$$

for any quantum state ρ . ${}_S^2 \hat{\mathbf{D}}^{pqrs}$ is the shadow estimator of the 2-RDM, and the subscript S differentiates it from, e.g., the analytic 2-RDM ${}^2\mathbf{D}^{pqrs}$ of a quantum state in the following.

The single-shot estimator for the expectation value of the 2-RDM for a measurement string b in the computational basis under the basis transformation generated by u is given by [11]

$${}_S^2 \hat{\mathbf{D}}^{pqrs} = \langle r, s | {}^2U(v_b u) E_{\eta, k=2} {}^2U^\dagger(v_b u) | p, q \rangle, \quad (5)$$

where the basis states $\langle r, s |$ and $| p, q \rangle$ are defined by $\langle r, s | = \langle 0 | a_s a_r$ and $| p, q \rangle = a_p^\dagger a_q^\dagger | 0 \rangle$ with the Fermi vacuum $| 0 \rangle$, indexing elements of the reduced density matrix. The matrix v_b implements a rotation of b into the ordered string $b_0 = \{11 \dots 00\} = \{1^{\otimes \eta}\} + \{0^{\otimes (N-\eta)}\}$. The fermionic rotations are combined for computational efficiency [11] ${}^kU(v_n) {}^kU_\eta(u) = {}^kU(v_b u)$. The operator $E_{\eta, k}$ is defined as

$$E_{\eta, k=2} = \sum_{p, q \in \{1, \dots, N\}} | p, q \rangle \langle p, q | \frac{\binom{\eta-s'}{2-s'} \binom{N-\eta+s'}{s'}}{(-1)^{2+s'} \binom{k}{s'}}, \quad (6)$$

with $s'(p, q)$ being the number of shared indices/occupied orbitals with the reference determinant/Hartree-Fock $| b_0 \rangle$:

$$s'(p, q) = \begin{cases} 2, & p \leq \eta \wedge q \leq \eta, \\ 1, & p \leq \eta \vee q \leq \eta, \\ 0, & p > \eta \wedge q > \eta. \end{cases}$$

The average shot variance of the 2-RDM estimator in Eq. (5) is then bounded by

$$\mathbb{E}_{pqrs} [\text{Var} [{}_S^2 \hat{\mathbf{D}}^{pqrs}]] \leq \binom{\eta}{2} \left(1 - \frac{\eta-2}{N} \right)^2 \left(\frac{1+N}{1+N-2} \right), \quad (7)$$

which scales as $\mathcal{O}(\eta^2)$. The estimator in Eq. (5) is only classically efficient for constant k [11], as is the case here for $k=2$.

III. v2rdm

Variational methods using the 2-RDM as a more compact stand-in for the wave function, as it contains all the information to calculate the energy, molecular forces, dipole moments, etc., were developed at the beginning of the 1950s [24]. Originally, obtaining energies that were too low [25], as unrestricted optimization results in 2-RDMs

not stemming from a physical wave function, was violating the variational principle. Once imbued with the concept of N -representability conditions, restricting the space of obtainable 1- and 2-RDMs in optimization, conditions in an ever-increasing order were developed to guide variational methods to obtain physical results when traversing the space of 2-RDMs, by e.g., a semidefinite program [26]. Recent high-performance implementations using graphics processing units of these methods have led to the capability to treat larger-scale systems, e.g., 3,7-circumjacent in a (64e, 64o) active space [16]. A brief treatment of the theory behind the workings of the algorithm will follow, but there are further resources [16] for a more in-depth introduction to the topic.

In general, v2rdm is based on a convex optimization in the form of a semidefinite program. The merit/cost function is formulated as a vector $\mathbf{c}$, containing the coefficients to score a solution vector $\mathbf{x}$ over an inner product $\mathbf{c}^T \mathbf{x}$ to be minimized. We use the energy under the Hamiltonian $E[\mathbf{D}]$ as the cost function. We compare different choices of cost functions in Appendix F. The N -representability conditions are formulated as a system of linear equations, containing the set of linear equations embedded into the rows of a matrix $\mathbf{A}$ and their respective solutions as a solution vector $\mathbf{b}$, while the entries of $\mathbf{x}$ are cast into a block diagonal form in a matrix $M(\mathbf{x})$, where the matrix and therefore the individual blocks are to be kept positive semidefinite during optimization. The full set of equations is summarized as

$$\begin{aligned} \min_{\mathbf{x}} \quad & \mathbf{c}^T \mathbf{x}, \\ \text{such that} \quad & \mathbf{A}\mathbf{x} = \mathbf{b}, \\ \text{and } M(\mathbf{x}) \quad & \succeq 0. \end{aligned} \quad (8)$$

In our case, the individual elements of $\mathbf{x}$ are associated with the individual entries of the 1-RDM ${}^1\mathbf{D}$ and 2-RDM ${}^2\mathbf{D}$, and further derived RDM-like quantities like the hole-particle 1-RDM ${}^1\mathbf{Q}^{pq} = \langle \psi | a_p a_q^\dagger | \psi \rangle$ and similar quantities derived from the 2-RDM, which appear in more elaborate N -representability conditions [15]. These are often referred to as the Q and G conditions, which come at an increased computational cost but lead to higher precision in the end result. Further constraints using the 3-RDM entries have been developed but will not be covered here. The block diagonal structure defined by $M(\mathbf{x})$ takes the form of

$$M(\mathbf{x}) = \begin{pmatrix} {}^1\mathbf{D} & 0 & 0 & 0 & 0 & 0 \\ 0 & {}^1\mathbf{Q} & 0 & 0 & 0 & 0 \\ 0 & 0 & {}^2\mathbf{D} & 0 & 0 & 0 \\ 0 & 0 & 0 & {}^2\mathbf{Q} & 0 & 0 \\ 0 & 0 & 0 & 0 & {}^2\mathbf{G} & 0 \\ 0 & 0 & 0 & 0 & 0 & \dots \end{pmatrix} \succeq 0. \quad (9)$$

As there is an exponential amount [27] of N -representability conditions, enforcing all of them is computationally unfeasible, and one usually truncates the conditions at conditions of second or third order. At the second order, there are efficient algorithms in solving the resulting semidefinite program scaling as $\mathcal{O}(N^4)$ in memory and $\mathcal{O}(N^6)$ in compute, e.g., a boundary-point method for semidefinite programming (BDSDP) [28–30], and a matrix-factorization-based procedure, reduced-rank semidefinite programming (RRSDP) [31]. In total, the

N -representability conditions usually considered in v2rdm can be summarized in the following way [15]:

$$\min_{{}^2\mathbf{D}} E[\mathbf{D}], \quad (10)$$

$$\text{such that } {}^2\mathbf{D} \succeq 0, \quad (11)$$

$${}^2\mathbf{Q} \succeq 0, \quad (12)$$

$${}^2\mathbf{G} \succeq 0, \quad (13)$$

$$\text{Tr}({}^2\mathbf{D}) = \eta(\eta - 1), \quad (14)$$

$${}^2\mathbf{Q} = f_Q({}^2\mathbf{D}), \quad (15)$$

$${}^2\mathbf{G} = f_G({}^2\mathbf{D}). \quad (16)$$

f_G and f_Q are linear mappings between ${}^2\mathbf{G}$, ${}^2\mathbf{Q}$, and ${}^2\mathbf{D}$. We use a spin-adapted v2rdm code and also enforce spin-adapted constraints from the classical shadow; for further details when adapting for spin, consult Appendix A. To these traditional constraints, we now add additional constraints stemming from the classical shadow measurements on a quantum computer.

IV. SHADOW CONSTRAINTS

We consider two different sets of constraints stemming from the output of a quantum computer, further constraining the optimization problem. The first way of constraining the optimization problem relies on the classical shadow estimator in Eq. (5), which bounds the variance of the average 2-RDM element in Eq. (7). As the variance properties of the classical shadow protocol rely on using the estimator, this is optimal in terms of quantum measurements. We refer to this constraint as constraint (1) from here. We build an ϵ ball around all the elements of the true 2-RDM ${}^2\mathbf{D}^{pqrs}$ from the estimator ${}_S\hat{\mathbf{D}}^{pqrs}$ in Eq. (5) and estimating the radius of the ball ϵ_1 by means of the upper bounds of the variance from Eq. (7):

$${}_S\hat{\mathbf{D}}^{pqrs} - \epsilon_1 \leq {}^2\mathbf{D}^{pqrs} \leq {}_S\hat{\mathbf{D}}^{pqrs} + \epsilon_1. \quad (17)$$

This results in $2N^4$ additional constraints to the optimization, which are of constant size. These can further be divided into restrictions on each spin sector as discussed in Appendix A. In general, the classical resources needed to calculate the estimator stemming from the classical shadow in Eq. (5) scale as $\mathcal{O}(N^5)$ in the classical postprocessing compared to the scaling $\mathcal{O}(N^6)$ of the semidefinite program. The classical optimization, therefore, adds a linear multiplicative overhead in classical computational cost while potentially reducing the number of quantum resources/shots required to reach a given precision. As the postmeasurement classical processing of the classical shadow comes from evaluating the single-shot estimator for each individual shot [32], reducing the necessary shots to reach a given precision can save computational cost from the preprocessing step, alleviating increased cost in the second step depending on the factor of improvement. There exists a more efficient way of estimating from a classical shadow using Pfaffians [10,11] not described here, which could bring this cost down even more.

The second constraint, referred to as constraint (2), introduced in Ref. [15], is based on the diagonal part of the

2-RDM, measured in the computational basis under m random single-particle rotations U_m as defined in Eq. (3):

$$\hat{S}_m^{pq} = \langle \Psi | U_m^\dagger a_p^\dagger a_q^\dagger a_q a_p U_m | \Psi \rangle, \quad (18)$$

with $|\Psi\rangle\langle\Psi| = \rho$. This constraint also relies on randomized measurements of the quantum state similar to (1) but not using any of the improved estimators provided by the classical shadow protocol. This leads to $2mN^2$ additional constraints, with N^2 entries. The factor of 2 follows from constraining both sides of the inequality (see Appendix B for more details):

$$\hat{S}_m^{pq} - \epsilon_2 \leq X_m^{pq} \leq \hat{S}_m^{pq} + \epsilon_2, \quad (19)$$

with

$$X_m^{pq} = ((U_m \otimes U_m)^2 D (U_m \otimes U_m)^T)^{pq}, \quad (20)$$

where ϵ_2 marks the error on each element due to shot noise and device error. We only investigate behavior under shot noise considerations in this article, but as device error can also be modeled under similar Gaussian perturbations to the correct RDM, one could further study behavior under noisy conditions [13,15].

V. NUMERICS

In this section, we numerically investigate the performance of the proposed optimization of the 2-RDM when using the two different constraints [Eqs. (17) and (19), labeled (1) and (2)] stemming from the classical shadow estimation in comparison with the estimate stemming from a vanilla v2rdm estimation without any additional constraints and an estimate stemming from a classical shadow estimation without any v2rdm postoptimization of the estimate [Eq. (5)]. To allow for comparison between all the approaches, we evaluate them at equal measurement shot budgets. The shot budget comparability is achieved in the following way: For more efficient numerical considerations, the estimate from a classical shadow for (1) is simulated by adding Gaussian noise to the 2-RDM from a full configuration interaction (FCI) calculation where the strength of the perturbation on each element is calculated by the bound of the shot variance of each element from Eq. (7). This is converted into a standard deviation by dividing by the number of shots taken at the respective point. Previous work [7,32] has shown this to be an accurate noise model for the estimate of the classical shadow. As the constraint for (2) in Eq. (19) relies on measuring the diagonal elements of the 2-RDM under random basis rotations, we consider the average variance of one of these diagonal elements. For, e.g., the ${}^2D_{\alpha\beta}$, the expectation value of a diagonal element of the 2-RDM $\mathbb{E}[\hat{S}_{pq}]$ can be calculated by

$$\mathbb{E}[\hat{S}_{pq}] = (N_\alpha N_\beta)^2 / \eta^2, \quad \text{Var}[\hat{S}_{pq}] = 1 - \mathbb{E}[\hat{S}_{pq}]^2, \quad (21)$$

where N_α, N_β are the number of α, β electrons and $\eta = N_\alpha + N_\beta$. The respective variance $\text{Var}[S_{pq}]$ follows from the observable mapping down to a Pauli observable (Z) when measuring in the computational basis after the random basis transformation. To allow for equal shots for both ways of constraining the problem, the number of basis m sampled when constraining with Eq. (19) is taken to fulfill

$$\mathbb{E}_{pqrs}[\text{Var}[\hat{S}_{pqrs}]] \approx m \text{Var}[\hat{S}_{pq}] \quad (22)$$

rounded up to the closest integer m . As the bound of the shot variance given in Eq. (7) is also only bounding the *average* shot variance of an entry of the 2-RDM, we deem this a fair comparison.

One way to view the two different constraints is as follows: Measuring the diagonal elements, one *oversamples* measurements under the same basis rotation, so from a measurement perspective, constraint (1) can be viewed as the limit of (2) when sampling in a basis for every shot and using the estimator in Eq. (5), which resembles Eq. (19) projected into the two-particle subspace. As changing the basis requires reconfiguring the quantum circuit, this is usually associated with increased cost, and *oversampling* has been explored in Ref. [7] for potentially reducing cost in shadow estimation when accounting for this cost. We do not include this potentially increased cost here, but it can also be accounted for when constraining in (1) by increasing the bound of the variance from Eq. (7). Covariances are not considered here, as previous work has shown numerical evidence that no covariances are present when estimating orthogonal elements in classical shadow estimation using matchgates [7,32].

Our implementation extends an open source version of a v2rdm code using libsdp, providing efficient solvers for the semidefinite program [16,27] with the constraints mentioned above. We further use Openfermion [33] to build the correct Hilbert space representation of our basis rotations in the correct subspace to build our constraints and simulate measurements on a quantum computer. Chemistry Hamiltonians and active space calculation are done by PySCF [34]. All simulations are run on a MacBook M1 Pro.

We use our test systems, N_2 in correlation-consistent polarized valence double zeta (cc-pvdz) basis in an (10, 8) active space and benzene in the Slater-type orbital basis STO-3G in a (16, 16) active space to investigate how system size affects the estimation. As previously mentioned, we are mostly concerned with improving the estimate of the 2-RDM from the shadow estimation scheme. Improvements in energy estimation, we deem less interesting as one has access to more efficient ways of measuring the energy, e.g., double factorization [35,36] or tensor hypercontraction [37]. Constructing the Hilbert-space unitaries needed for (2) becomes fairly costly at large system sizes, and consequently, we only investigate constraint (1) for the largest of the test systems, benzene.

For the smaller active space (10, 8) for N_2 , looking at the energy, we find shadow constraint (1) outperforming shadow constraint (2) over the range of the shot interval investigated (Fig. 1, top). For the estimate of the 2-RDM compared to the true 2-RDM (Fig. 1, bottom), we also find shadow constraint (1) outperforming constraint (2) consistently. Above 10^6 shots, shadow constraint (1) improves over the plain v2rdm implementation, whereas the unadapted shadow estimation only improves above 10^7 over the purely classical estimate. Constraint (1) improves over the shadow estimate up until 2×10^8 shots, where the shadow estimator overtakes the combined estimate, reaching a reduction of up to a factor of 15 required shots to reach the same precision. The threshold at 2×10^8 can be explained by the energy acting as the cost function in optimization: At some point, the optimizer biases the estimate by improving the entries of the 2-RDM with high weight from the Hamiltonian at the cost of the other entries,




FIG. 1. Top: Energy error of the estimate of the ground-state energy for N_2 in dependence on the number of shots spent on the quantum shadow constraints included in the optimization. Bottom: 2-RDM error in the Frobenius norm toward the true 2-RDM stemming from an FCI calculation of N_2 in dependence on the number of shots spent on the quantum shadow constraints included in the optimization. Inset quantifies the factor of improvement over the estimate coming just from the classical shadow at the same level of error (red). The dot in the inset corresponds to the dashed line in the outer figure as a guide to the eye and clarification.

losing precision over the unbiased estimate of the classical shadow estimation performed on the true ground state (see Appendix E). Shadow constraint (2) does not improve over v2rdm without constraints in the low shot regime or the estimate from classical shadow estimation at the 10^7 shots threshold.




FIG. 2. 2-RDM error in the Frobenius norm toward the true 2-RDM stemming from an FCI calculation of benzene in a (16, 16) active space in dependence of the number of shots spent on the quantum shadow constraints included in the optimization. Inset quantifies the factor of improvement over the estimate coming just from the classical shadow at the same level of error (red). The dot in the inset corresponds to the dashed line in the outer figure as a guide to the eye and clarification.

In Appendix G, adding more information from a fault-tolerant calculation, e.g., the estimate of the ground-state energy is investigated.

For the large active space (16, 16) for benzene (Fig. 2), we find a qualitatively very similar looking relationship between plain v2rdm, classical shadow estimation, and v2rdm with shadow constraint (1): At an error that corresponds roughly with 10^{-3} , the shadow estimation beats the purely classical estimate. v2rdm with shadow constraint (1) improves over both the other estimates for roughly 2 orders of magnitude of shot budgets before the bias introduced into the estimate loses over the unbiased shadow estimation technique. Improvement factors also seem to be a bit lower in this scenario, only reaching up to a factor of 5.

VI. DISCUSSION

To conclude, our simulations show that v2rdm can help classical shadow estimation to increase precision for a given shot budget in certain regimes. When accounting for shot noise, we find that constraint (2) of previous work [15] does not improve the 2-RDM estimate over either the only classical optimization or quantum-only estimator at a given shot budget. Only when constraining the optimization with the improved constraint (1), we find up to a factor of 15 in improvement for N_2 , although the factor of improvement is largest when being close to the result of the purely classical

estimate. This highlights that the power of the classical shadow techniques stems from the efficient estimators that come with guarantees of precision and they should be included in the procedure and fully randomizing the measurement scheme instead of oversampling a limited set of basis rotations. The improvement we find for constraint (2) seems to be reduced for larger systems. For high precision estimates, the bias introduced seems to lose out over the unbiased estimate from shadow estimation. Nonetheless, when quantum resources are scarce in near- and midterm shadow estimations trying to extract 2-RDM quantities from a state to extract, e.g., forces and other molecular properties beyond ground-state energies, classically optimizing the estimate from a classical shadow can lead to increased precision at a fixed quantum cost. Furthermore, when comparing different choices of cost function, minimizing the energy under the Hamiltonian still seems to be the leading choice, even when adding the exact ground-state energy as a constraint in the other cases (Appendixes F and G). Although the window where v2rdm actually improves the estimate is limited, for estimations where high precision is necessary, e.g., chemistry applications, this does not solve the cost of extracting a 2-RDM from a quantum device. As we have investigated scaling of epsilon under the standard quantum limit $\frac{1}{\epsilon^2}$, the cost savings would become much more dramatic for recently proposed procedures [38] with better scaling in N at the cost of increased scaling in the error ϵ . Furthermore, there is also potential to use this approach as an error mitigation technique beyond shot noise when allowing for a bias in the error model [13,39].

As the optimization under the energy as a cost function seems to introduce a bias limiting the applicability window where the combined methods improve over pure shadow estimation, changing the cost function to, e.g., the 1-norm of the 2-RDM [13,15] might be an avenue in further improvement of the scheme presented here.

ACKNOWLEDGMENTS

We thank Nicholas Rubin and Ryan Babbush for discussions on fermionic marginal constraints. Furthermore, we thank Christian Gogolin and Maximilian Scheurer for input on shadow estimation. We thank Eugene DePrince III and Brecht Verstichel for discussions about the v2rdm method.

DATA AVAILABILITY

The data that support the findings of this article are not publicly available. The data are available from the authors upon reasonable request.

APPENDIX A: SPIN-RESTRICTED v2rdm

As we work in a spin-adapted setting, one can further find a subblock structure of the positive-semidefinite matrices. We decompose the one-particle density matrix ${}^1\mathbf{D}^{pq} = \langle \psi | a_p^\dagger a_q | \psi \rangle$ and its derivative ${}^1\mathbf{Q}^{pq} = \langle \psi | a_p a_q^\dagger | \psi \rangle$ into spin blocks in the following way, as discussed in further depth in Ref. [13]:

$${}^1\mathbf{D} = \begin{pmatrix} {}^1\mathbf{D}_\alpha & 0 \\ 0 & {}^1\mathbf{D}_\beta \end{pmatrix} \quad (\text{A1})$$

and $\mathbf{Q}$:

$${}^1\mathbf{Q} = \begin{pmatrix} {}^1\mathbf{Q}_\alpha & 0 \\ 0 & {}^1\mathbf{Q}_\beta \end{pmatrix}. \quad (\text{A2})$$

The two-particle conditions of ${}^2\mathbf{D}^{pqrs} = \langle \psi | a_p^\dagger a_q^\dagger a_s a_r | \psi \rangle$ and derived quantities ${}^2\mathbf{Q}^{pqrs} = \langle \psi | a_p^\dagger a_q a_s^\dagger a_r | \psi \rangle$ and ${}^2\mathbf{Q}^{pqrs} = \langle \psi | a_p a_q a_s^\dagger a_r^\dagger | \psi \rangle$ can also further be block diagonalized into spin sectors. For compactness sake, we leave out the upper indices on the diagonal ${}^2\mathbf{D}_{\alpha\alpha} \equiv {}^2\mathbf{D}_{\alpha\alpha}^{\alpha\alpha}$ and get four spin blocks for the 2-RDM

$${}^2\mathbf{D} = \begin{pmatrix} {}^2\mathbf{D}_{\alpha\alpha} & 0 & 0 & 0 \\ 0 & {}^1\mathbf{D}_{\beta\beta} & 0 & 0 \\ 0 & 0 & {}^2\mathbf{D}_{\alpha\beta} & 0 \\ 0 & 0 & 0 & {}^2\mathbf{D}_{\beta\alpha} \end{pmatrix} \quad (\text{A3})$$

and a slightly more elaborate structure of G :

$${}^2\mathbf{G} = \begin{pmatrix} {}^2\mathbf{G}_{\alpha,\alpha} & {}^2\mathbf{G}_{\alpha,\alpha}^{\beta,\beta} & 0 & 0 \\ {}^2\mathbf{G}_{\beta,\beta}^{\alpha,\alpha} & {}^2\mathbf{G}_{\beta,\beta} & 0 & 0 \\ 0 & 0 & {}^2\mathbf{G}_{\alpha,\beta} & 0 \\ 0 & 0 & 0 & {}^2\mathbf{G}_{\beta,\alpha} \end{pmatrix}, \quad (\text{A4})$$

whereas $\mathbf{Q}$ has the same structure as G .

APPENDIX B: INEQUALITIES IN SEMIDEFINITE PROGRAMMING

As we are constraining our semidefinite program now with estimates associated with a certain error ϵ coming from the statistical nature of the estimate, we have to be able to handle inequalities when constraining the optimization problem. To accommodate inequalities, one converts these inequalities into equalities by using a set of ancillary variables $\varphi_0, \dots, \varphi_i$ that are added to the optimization in the form of individual 1 by 1 blocks:

$$M(\mathbf{x}) = \begin{pmatrix} {}^1\mathbf{D} & 0 & \dots \\ 0 & \ddots & \ddots \\ \vdots & \ddots & {}^2\mathbf{D} & \ddots \\ & & \ddots & \ddots & \ddots \\ & & & \ddots & \varphi_0 & \ddots & \vdots \\ & & & & \ddots & \ddots & 0 \\ & & & & & \dots & 0 & \varphi_i \end{pmatrix} \succeq 0. \quad (\text{B1})$$

The j th inequality from Eq. (17) constraining the problem

$${}^2\hat{\mathbf{D}}^{pqrs} - \epsilon \leq {}^2\mathbf{D}^{pqrs} \leq {}^2\hat{\mathbf{D}}^{pqrs} + \epsilon \quad (\text{B2})$$

can then be converted into two equalities by adding an ancillary variable to a side of each inequality to equalize both sides

$${}^2\hat{\mathbf{D}}^{pqrs} - \epsilon = {}^2\mathbf{D}^{pqrs} - \varphi_{2j}, \quad (\text{B3})$$

$${}^2\mathbf{D}^{pqrs} + \varphi_{2j+1} = {}^2\hat{\mathbf{D}}^{pqrs} + \epsilon, \quad (\text{B4})$$

where j is an index keeping track of labeling individual sets of p, q, r, s , assigning it its own pair of ancillary variables. This holds as long as $\varphi_i > 0$, which is guaranteed by keeping the individual ancillary variables φ_i in an individual (1×1) block of $M(x)$ during the optimization, and the maximum error on

each element can be upper bounded by Eq. (7) for constraint (1).

In the spin-adapted way of constraining the semidefinite program, our condition from Eq. (17) becomes a restriction on the individual spin-full 2-RDMs

$$\frac{2}{S}\hat{\mathbf{D}}_{\alpha\alpha}^{pqrs} - \epsilon \leq \frac{2}{S}\mathbf{D}_{\alpha\alpha}^{pqrs} \leq \frac{2}{S}\hat{\mathbf{D}}_{\alpha\alpha}^{pqrs} + \epsilon, \quad (\text{B5})$$

$$\frac{2}{S}\hat{\mathbf{D}}_{\alpha\beta}^{pqrs} - \epsilon \leq \frac{2}{S}\mathbf{D}_{\alpha\beta}^{pqrs} \leq \frac{2}{S}\hat{\mathbf{D}}_{\alpha\beta}^{pqrs} + \epsilon, \quad (\text{B6})$$

$$\frac{2}{S}\hat{\mathbf{D}}_{\beta\beta}^{pqrs} - \epsilon \leq \frac{2}{S}\mathbf{D}_{\beta\beta}^{pqrs} \leq \frac{2}{S}\hat{\mathbf{D}}_{\beta\beta}^{pqrs} + \epsilon. \quad (\text{B7})$$

Again, each of these constraints can be converted into two equalities by adding an ancillary variable to a side of the inequality, here shown for one of the (α, α) constraints (B5):

$$\frac{2}{S}\hat{\mathbf{D}}_{\alpha\alpha}^{pqrs} - \epsilon = \frac{2}{S}\mathbf{D}_{\alpha\alpha}^{pqrs} - \varphi_{2i,\alpha\alpha}, \quad (\text{B8})$$

$$\frac{2}{S}\mathbf{D}_{\alpha\alpha}^{pqrs} + \varphi_{2i+1,\alpha\alpha} = \frac{2}{S}\hat{\mathbf{D}}_{\alpha\alpha}^{pqrs} + \epsilon. \quad (\text{B9})$$

More details on how to extract spin-full RDMs from the classical shadow are found in Appendix C.

APPENDIX C: SPIN RESTRICTION OF CLASSICAL SHADOWS

One can further adapt the estimator in Eq. (5) and the shadow protocol to respect another symmetry apart from particle number—spin. One can restrict the Haar-random unitary generating the basis rotation u_{pq} into a block-diagonal form [40],

$$u_{pq} = \begin{pmatrix} u_{pq,\alpha} & 0 \\ 0 & u_{pq,\beta} \end{pmatrix}, \quad (\text{C1})$$

which allows for a more efficient way of implementing the unitary in terms of Givens rotations into two local Givens circuits reducing the circuit depth needed to implement the measurement rotation further by a factor of 2. As the classical postprocessing time scales with the size of the Hilbert subspace, one can carefully restrict the estimator to a smaller subspace

$$E_{\eta,2}^{\alpha,\beta} = \sum_{p,q \in \{1, \dots, N_{\text{orb}}\}} |p^\alpha, q^\beta\rangle \langle p^\alpha, q^\beta| \frac{\binom{\eta-s'}{k-s'} \binom{N-\eta+s'}{s'}}{(-1)^{k+s'} \binom{k}{s'}}, \quad (\text{C2})$$

likewise for (α, α) , (β, α) , and (β, β) to only index the α/β part of the p/q spatial orbital. This reduces the size of the individual subspaces by a factor of 2 and the number of entries by factor of 4.

APPENDIX D: ADDITIONAL NUMERICS

To verify our implementation of shadow constraint (2), we also run simulations in the scenario considered most of the time in Ref. [15] at $\epsilon_2 = 0$. We study H₄ in STO-3G, the energy convergence in Fig. 3 in dependence of the number of basis sampled m , also reaching exact agreement at $m = 12$ basis used for the additional constraints during optimization at infinite precision, corresponding to taking an infinite number of shots at each basis. So, we deem our implementation of the shadow constraint comparable to the one studied in previous work.




FIG. 3. Energy convergence for H₄ for the optimization under shadow constraint (2) in dependence of number of basis m used for additional constraints during optimization.

APPENDIX E: INTRODUCED BIAS IN OPTIMIZATION

As our numerics show, when including a high precision estimate of a classical shadow with a tight ϵ , the optimization actually increases the distance to the true 2-RDM during optimization (Fig. 3). Here, we investigate the mechanism in how the bias is introduced. We look at the elements of the 2-RDM after the semidefinite optimization at different precisions of the ingoing shadow data. We linearize the matrix and sort it by the magnitude of the associated coefficients (see Fig. 4, top) in the Hamiltonian that constitutes the cost function used in the optimization.

We plot the difference of each element to the estimate from the classical shadow estimator used to generate the additional constraints. As the shadow data are generated by applying




FIG. 4. Top: Absolute value of the two-body coefficients of the Hamiltonian/two electron integrals for H₄ in STO-3G sorted descending. Bottom: Difference between the shadow estimate at a given epsilon used to build the shadow constraints and the 2-RDM estimate after the semidefinite optimization averaged over five runs sorted by the ordering of the upper plot.

Gaussian noise to the correct 2-RDM, each element will be on average equally far away from its true value and therefore should be moved similarly during optimization. At low precision (Fig. 4, bottom, blue), we can see that elements associated with big and small coefficients in the Hamiltonian get changed in the optimization on similar magnitudes; therefore, no or at least a very small bias is introduced compared to the variance of the estimator. At high precision (Fig. 4, bottom, red), one can clearly identify the point at which the Hamiltonian coefficients decrease by 6 orders of magnitude in the bottom plot. To decrease the energy that constitutes the cost function, terms with high weight get moved an order of magnitude more during the optimization than lower weight terms introducing a bias to the result, as in an unbiased optimization we expect each element to be moved similarly during optimization.

APPENDIX F: CHOICE OF COST FUNCTION

When variationally optimizing the 2-RDM with regard to N -representability constraints, most implementations [15,27,30] minimize the energy $\min_{\mathbf{D}} E[\mathbf{D}]$ with regard to the Hamiltonian as a cost function [27,30]. Other forms of cost functions have been used very successfully in other applications of semidefinite programming, e.g., when completing matrices with only partially known entries [41] and recovering matrices from noisy estimates of their entries [42]. This is based on the finding that low-rank approximations of these matrices are on average a good representation of the unknown matrix. As minimizing the rank is NP-hard [43], the first relaxation to a convex manifold being the trace or nuclear norm $\|\cdot\|_*$ is used as a heuristic in optimization [43,44]. The Frobenius norm $\|\cdot\|_F$ also serves as a possible cost function in these applications [15]. For a (positive-semidefinite) matrix A , the norms are defined as

$$\|A\|_* = \text{Tr}(\sqrt{A^*A}), \quad (\text{F1})$$

$$\|A\|_F = \sqrt{\text{Tr}(A^*A)}. \quad (\text{F2})$$

These type of cost functions have also been applied in quantum chemistry settings, as 2-RDMs of ground states are known to be approximately low rank [42,45,46] and have been investigated in improving readout when sampling in the Pauli basis [42], finding that in the noisy readout setting it was optimal to sample all terms and the optimization mostly acting as a low-rank filter to the estimation. As the shadow estimation used here estimates all Fermionic terms with equal error, optimizing for nuclear norm also seems as a reasonable choice as a cost function.

We investigate the following different cost functions in our setting:

$$\min_{\mathbf{D}} E[\mathbf{D}], \quad (\text{F3})$$

$$\min_{\mathbf{D}} E[\mathbf{D}] - \|\mathbf{D}\|_*, \quad (\text{F4})$$

$$\min_{\mathbf{D}} \|\mathbf{D}\|_*, \quad (\text{F5})$$

$$\min_{\mathbf{D}} \|\mathbf{D}\|_F. \quad (\text{F6})$$

There are a total of four different cost functions, we consider the following:




FIG. 5. Comparison of different choices of cost functions described in Eqs. (F3)–(F6) against the plain error on the shadow estimator entering the constraints. The constraints otherwise are kept the same; points are averaged over multiple runs with different errors on the shadow constraints.

- (1) The energy under the Hamiltonian (F3).
- (2) The energy under the Hamiltonian regularized by the nuclear norm (Schatten 1-norm or trace norm) $\|\cdot\|_1$ of the 2-RDM, therefore penalizing solutions with higher nuclear norm (F4).
- (3) The nuclear norm $\|\cdot\|_1$ of the 2-RDM (F5).
- (4) The Frobenius norm $\|\cdot\|_F$ of the 2-RDM (F6).

The constraints are the same for all the cost functions. The optimization is implemented by using the package CVXPY [47] and plotted in Fig. 5. The test system is the (10, 8) active space for N_2 in the cc-pvdz basis set used in the main text as a testing system. The 2-RDM is cast into a matrix form via a superindex summing p, q and r, s as is already done when setting up the constraints in v2rdm.

We do not find a significant overall difference between the different cost functions. In the noisy, fairly unconstrained setting that extrapolates to the plain v2rdm algorithm the cost functions using energy and the regularized energy clearly dominate, as no real information about the Hamiltonian or its ground state enters the optimization at a certain noise level. In the fairly constrained setting, the cost functions only use rank seem to introduce less of a bias over the shadow estimator but still ending up less accurate as the shadow estimator in the high precision regime.

So, at least from this small numerical experiment it does not seem that extracting approximate low-rank 2-RDMs from a noisy classical shadow estimation seems to improve over the optimization under energy or the pure shadow estimator.

1. Implementing norm cost functions in semidefinite programming

As on first glance, it is quite unintuitive how to convert cost functions including the nuclear or Frobenius into the linear constraints needed for a semidefinite program; we quickly highlight the modifications needed to transform them into the appropriate format. When using the nuclear norm as a minimization target, one arrives at the following semidefinite program:

$$\begin{aligned} \min_A \quad & \|A\|_*, \\ \text{subject to} \quad & \mathcal{C}(A) = b, \end{aligned} \quad (\text{F7})$$

where $\mathcal{C}$ represents linear constraints. This problem can be reformulated [48] as

$$\begin{aligned} \min_{A,U,V} \quad & \frac{1}{2}(\text{tr}(U) + \text{tr}(V)), \\ \text{subject to} \quad & \begin{bmatrix} U & A \\ A^\dagger & V \end{bmatrix} \succeq 0, \\ & \mathcal{C}(A) = b, \end{aligned} \quad (\text{F8})$$

where $U, V \in \mathbb{H}^n$ are auxiliary semidefinite variables.

A very similar thing can be done when minimizing under the Frobenius norm

$$\begin{aligned} \min_A \quad & \|A\|_F, \\ \text{subject to} \quad & \mathcal{C}(A) = b. \end{aligned} \quad (\text{F9})$$

The reformulation in this case using the Schur complement [49,50] becomes

$$\begin{aligned} \min_{A,W,t} \quad & t, \\ \text{subject to} \quad & \begin{bmatrix} W & A \\ A^\dagger & I \end{bmatrix} \succeq 0, \\ & \text{Tr}(W) \leq t, \\ & \mathcal{C}(A) = b, \end{aligned} \quad (\text{F10})$$

where the semidefinite constraint together with the trace bound encodes $\|A\|_F^2 \leq t$. By the Schur complement lemma (see Ref. [49], §4.6.3), the matrix inequality implies $W \succeq AA^\dagger$, so that $\text{Tr}(W) \geq \text{Tr}(AA^\dagger) = \|A\|_F^2$. As these are regularly showing up in semidefinite programs, these reformulations are already implemented for the practising researcher using semidefinite programs in open source semidefinite program solvers like CVXPY [47] and can be used without deriving these reformulations.

One might also wonder whether constructing constraints based on these norms could improve the optimization further. A natural approach would be to replace the elementwise constraints with a single norm-based constraint of the form $\|\hat{\mathbf{D}} - \hat{\mathbf{D}}\| \leq \epsilon$, which can also be formulated within a semidefinite program framework using the reformulations above. Under the Gaussian noise model used in our numerical simulations, concentration results from random matrix theory could in principle be used to set ϵ : For a $d \times d$ Hermitian error matrix with independent and identically distributed Gaussian entries of variance σ^2 , the Frobenius norm concentrates around σd and the operator norm around $2\sigma\sqrt{d}$ [51,52]. However, we

do not think that our assumption of the Gaussian distribution of noise holds in practice; therefore, we do not find rigorous ways of including norm-based constraints to our optimization, though this remains a promising direction for future work.

APPENDIX G: CONSTRAINING WITH GROUND-STATE ENERGY

In the fault-tolerant regime, quantum phase estimation provides access to the ground-state energy under controlled approximation under certain conditions at reasonable cost. Although providing access to the ground-state energy, reduced density matrices are still prohibitively expensive when looking to estimate further molecular observables like, e.g., forces for geometric optimization. Therefore, one can pose the question whether constraining the energy by the ground-state energy in a semidefinite program allows for a more accurate estimation.

To include the estimate of the ground-state energy in our semidefinite program, we add the following constraint to the optimization:

$$E_{\text{GS}} - \epsilon_{\text{GS}} \leq \text{Tr}(H^2 \mathbf{D}) \leq E_{\text{GS}} + \epsilon_{\text{GS}}, \quad (\text{G1})$$

where E_{GS} is the ground-state energy (calculated by FCI in PySCF) and $\epsilon_{\text{GS}} = 10^{-3}$ hartree at chemical accuracy, the precision a quantum phase estimation estimate would aim to achieve. The test system is again the (10, 8) active space for N_2 in the cc-pvdz basis set used in the main text. Results are plotted in Fig. 6.

In our limited data set, we do not find an improvement over not including the ground-state energy in our optimization. Adding this constraint does not seem to be restricting our allowed solution space in a meaningful way.




FIG. 6. Comparison of optimizing the nuclear norm of cost functions described in Eq. (F5) when omitting/adding the energy constraint from Eq. (G1) against the plain error on the shadow estimator. Points are averaged over multiple runs.

- [1] D. S. Abrams and S. Lloyd, Simulation of many-body Fermi systems on a universal quantum computer, *Phys. Rev. Lett.* **79**, 2586 (1997).
- [2] A. Aspuru-Guzik, A. D. Dutoi, P. J. Love, and M. Head-Gordon, Simulated quantum computation of molecular energies, *Science* **309**, 1704 (2005).
- [3] Y. Cao, J. Romero, J. P. Olson, M. Degroote, P. D. Johnson, M. Kieferová, I. D. Kivlichan, T. Menke, B. Peropadre, N. P. D. Sawaya, S. Sim, L. Veis, and A. Aspuru-Guzik, Quantum chemistry in the age of quantum computing, *Chem. Rev.* **119**, 10856 (2019).
- [4] S. McArdle, S. Endo, A. Aspuru-Guzik, S. C. Benjamin, and X. Yuan, Quantum computational chemistry, *Rev. Mod. Phys.* **92**, 015003 (2020).
- [5] J. Preskill, Quantum computing in the NISQ era and beyond, *Quantum* **2**, 79 (2018).
- [6] 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 (London)* **603**, 416 (2022).
- [7] M. Scheurer, G.-L. R. Anselmetti, O. Oumarou, C. Gogolin, and N. C. Rubin, Tailored and externally corrected coupled cluster with quantum inputs, *J. Chem. Theory Comput.* **20**, 5068 (2024).
- [8] H.-Y. Huang, R. Kueng, and J. Preskill, Predicting many properties of a quantum system from very few measurements, *Nat. Phys.* **16**, 1050 (2020).
- [9] A. Zhao, N. C. Rubin, and A. Miyake, Fermionic partial tomography via classical shadows, *Phys. Rev. Lett.* **127**, 110504 (2021).
- [10] K. Wan, M. Berta, and E. T. Campbell, Randomized quantum algorithm for statistical phase estimation, *Phys. Rev. Lett.* **129**, 030503 (2022).
- [11] G. H. Low, Classical shadows of fermions with particle number symmetry, *arXiv:2208.08964*.
- [12] H.-Y. Huang, M. Broughton, J. Cotler, S. Chen, J. Li, M. Mohseni, H. Neven, R. Babbush, R. Kueng, J. Preskill, and J. R. McClean, Quantum advantage in learning from experiments, *Science* **376**, 1182 (2022).
- [13] N. C. Rubin, R. Babbush, and J. McClean, Application of fermionic marginal constraints to hybrid quantum algorithms, *New J. Phys.* **20**, 053020 (2018).
- [14] F. Arute *et al.*, Hartree-Fock on a superconducting qubit quantum computer, *Science* **369**, 1084 (2020).
- [15] I. Avdic and D. A. Mazziotti, Fewer measurements from shadow tomography with N -representability conditions, *Phys. Rev. Lett.* **132**, 220802 (2024).
- [16] A. Eugene DePrince III, Variational determination of the two-electron reduced density matrix: A tutorial review, *WIREs Comput. Mol. Sci.* **14**, e1702 (2024).
- [17] B. Verstichel, H. van Aggelen, D. Van Neck, P. W. Ayers, and P. Bultinck, Variational density matrix optimization using semidefinite programming, *Comput. Phys. Commun.* **182**, 2025 (2011), Computer Physics Communications Special Edition for Conference on Computational Physics Trondheim, Norway, June 23–26, 2010.
- [18] S. Simon, R. Santagati, M. Degroote, N. Moll, M. Streif, and N. Wiebe, Improved precision scaling for simulating coupled quantum-classical dynamics, *PRX Quantum* **5**, 010343 (2024).
- [19] S. Simon, M. Degroote, N. Moll, R. Santagati, M. Streif, and N. Wiebe, Amplified amplitude estimation: Exploiting prior knowledge to improve estimates of expectation values, *arXiv:2402.14791*.
- [20] R. Santagati, A. Aspuru-Guzik, R. Babbush, M. Degroote, L. González, E. Kyoseva, N. Moll, M. Oppel, R. M. Parrish, N. C. Rubin, M. Streif, C. S. Tautermann, H. Weiss, N. Wiebe, and C. Utschig-Utschig, Drug design on quantum computers, *Nat. Phys.* **20**, 549 (2024).
- [21] T. E. O’Brien, M. Streif, N. C. Rubin, R. Santagati, Y. Su, W. J. Huggins, J. J. Goings, N. Moll, E. Kyoseva, M. Degroote, C. S. Tautermann, J. Lee, D. W. Berry, N. Wiebe, and R. Babbush, Efficient quantum computation of molecular forces and other energy gradients, *Phys. Rev. Res.* **4**, 043210 (2022).
- [22] M. Steudtner, S. Morley-Short, W. Pol, S. Sim, C. L. Cortes, M. Loipersberger, R. M. Parrish, M. Degroote, N. Moll, R. Santagati, and M. Streif, Fault-tolerant quantum computation of molecular observables, *Quantum* **7**, 1164 (2023).
- [23] S. Aaronson, Shadow tomography of quantum states, *SIAM J. Comput.* **49**, STOC18-368 (2020).
- [24] P.-O. Löwdin, Quantum theory of many-particle systems. I. Physical interpretations by means of density matrices, natural spin-orbitals, and convergence problems in the method of configurational interaction, *Phys. Rev.* **97**, 1474 (1955).
- [25] J. E. Mayer, Electron correlation, *Phys. Rev.* **100**, 1579 (1955).
- [26] 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.* **114**, 8282 (2001).
- [27] A. E. DePrince III, libSDP: A library of semidefinite programming solvers, 2021, <https://github.com/edeprince3/libsdp>.
- [28] J. Povh, F. Rendl, and A. Wiegeler, A boundary point method to solve semidefinite programs, *Computing* **78**, 277 (2006).
- [29] J. Malick, J. Povh, F. Rendl, and A. Wiegeler, Regularization methods for semidefinite programming, *SIAM J. Optim.* **20**, 336 (2009).
- [30] D. A. Mazziotti, Large-scale semidefinite programming for many-electron quantum mechanics, *Phys. Rev. Lett.* **106**, 083001 (2011).
- [31] D. A. Mazziotti, Realization of quantum chemistry without wave functions through first-order semidefinite programming, *Phys. Rev. Lett.* **93**, 213001 (2004).
- [32] M. Kiser, A. Schroeder, G.-L. R. Anselmetti, C. Kumar, N. Moll, M. Streif, and D. Vodola, Classical and quantum cost of measurement strategies for quantum-enhanced auxiliary field quantum Monte Carlo, *New J. Phys.* **26**, 033022 (2024).
- [33] J. R. McClean *et al.*, OpenFermion: The electronic structure package for quantum computers, *Quantum Sci. Technol.* **5**, 034014 (2020).
- [34] Q. Sun, T. C. Berkelbach, N. S. Blunt, G. H. Booth, S. Guo, Z. Li, J. Liu, J. McClain, E. R. Sayfutyarova, S. Sharma, S. Wouters, and G. K.-L. Chan, The Python-based simulations of chemistry framework (PySCF), *WIREs Comput. Mol. Sci.* **8**, e1340 (2018).
- [35] V. von Burg, G. H. Low, T. Häner, D. S. Steiger, M. Reiher, M. Roetteler, and M. Troyer, Quantum computing enhanced computational catalysis, *Phys. Rev. Res.* **3**, 033055 (2021).
- [36] O. Oumarou, M. Scheurer, R. M. Parrish, E. G. Hohenstein, and C. Gogolin, Accelerating quantum computations of chemistry

- through regularized compressed double factorization, [Quantum](#) **8**, 1371 (2024).
- [37] J. Lee, D. W. Berry, C. Gidney, W. J. Huggins, J. R. McClean, N. Wiebe, and R. Babbush, Even more efficient quantum computations of chemistry through tensor hypercontraction, [PRX Quantum](#) **2**, 030305 (2021).
- [38] R. King, D. Gosset, R. Kothari, and R. Babbush, Triply efficient shadow tomography, [PRX Quantum](#) **6**, 010336 (2025).
- [39] Z. Cai, R. Babbush, S. C. Benjamin, S. Endo, W. J. Huggins, Y. Li, J. R. McClean, and T. E. O'Brien, Quantum error mitigation, [Rev. Mod. Phys.](#) **95**, 045005 (2023).
- [40] A. Zhao and A. Miyake, Group-theoretic error mitigation enabled by classical shadows and symmetries, [npj Quantum Inf.](#) **10**, 57 (2024).
- [41] E. J. Candes and B. Recht, Exact low-rank matrix completion via convex optimization, in *2008 46th Annual Allerton Conference on Communication, Control, and Computing* (IEEE, Piscataway, NJ, 2008), pp. 806–812.
- [42] L. Peng, X. Zhang, and G. K.-L. Chan, Fermionic reduced density low-rank matrix completion, noise filtering, and measurement reduction in quantum simulations, [J. Chem. Theory Comput.](#) **19**, 9151 (2023).
- [43] Y. Chi, Low-rank matrix completion [lecture notes], [IEEE Signal Process. Mag.](#) **35**, 178 (2018).
- [44] D. Gross, Recovering low-rank matrices from few coefficients in any basis, [IEEE Trans. Inf. Theory](#) **57**, 1548 (2011).
- [45] C. A. Schwerdtfeger and D. A. Mazziotti, Low-rank spectral expansions of two electron excitations for the acceleration of quantum chemistry calculations, [J. Chem. Phys.](#) **137**, 244103 (2012).
- [46] G. Gidofalvi and D. A. Mazziotti, Multireference self-consistent-field energies without the many-electron wave function through a variational low-rank two-electron reduced-density-matrix method, [J. Chem. Phys.](#) **127**, 244105 (2007).
- [47] S. Diamond and S. Boyd, CVXPY: A Python-embedded modeling language for convex optimization, [J. Mach. Learn. Res.](#) **17**, 1 (2016).
- [48] B. Recht, M. Fazel, and P. A. Parrilo, Guaranteed minimum-rank solutions of linear matrix equations via nuclear norm minimization, [SIAM Rev.](#) **52**, 471 (2010).
- [49] S. Boyd and L. Vandenberghe, *Convex Optimization* (Cambridge University Press, Cambridge, England, 2004).
- [50] L. El Ghaoui and H. Le Bret, Robust solutions to least-squares problems with uncertain data, [SIAM J. Matrix Anal. Appl.](#) **18**, 1035 (1997).
- [51] R. Vershynin, Introduction to the non-asymptotic analysis of random matrices, in *Compressed Sensing: Theory and Applications* (Cambridge University Press, 2012).
- [52] G. W. Anderson, A. Guionnet, and O. Zeitouni, *An Introduction to Random Matrices* (Cambridge University Press, Cambridge, England, 2010).