

Efficient Estimation of Pauli Observables by Derandomization

Hsin-Yuan Huang^{1,2,*}, Richard Kueng³, and John Preskill^{1,2,4,5}

¹*Institute for Quantum Information and Matter, Caltech, Pasadena, California 91125, USA*

²*Department of Computing and Mathematical Sciences, Caltech, Pasadena, California 91125, USA*

³*Institute for Integrated Circuits, Johannes Kepler University Linz, A-4040, Austria*

⁴*Walter Burke Institute for Theoretical Physics, Caltech, Pasadena, California 91125, USA*

⁵*AWS Center for Quantum Computing, Pasadena, California 91125, USA*



(Received 19 March 2021; accepted 14 June 2021; published 16 July 2021)

We consider the problem of jointly estimating expectation values of many Pauli observables, a crucial subroutine in variational quantum algorithms. Starting with randomized measurements, we propose an efficient derandomization procedure that iteratively replaces random single-qubit measurements by fixed Pauli measurements; the resulting deterministic measurement procedure is guaranteed to perform at least as well as the randomized one. In particular, for estimating any L low-weight Pauli observables, a deterministic measurement on only of order $\log(L)$ copies of a quantum state suffices. In some cases, for example, when some of the Pauli observables have high weight, the derandomized procedure is substantially better than the randomized one. Specifically, numerical experiments highlight the advantages of our derandomized protocol over various previous methods for estimating the ground-state energies of small molecules.

DOI: [10.1103/PhysRevLett.127.030503](https://doi.org/10.1103/PhysRevLett.127.030503)

Introduction.—Noisy intermediate-scale quantum (NISQ) devices are becoming available [1]. Though less powerful than fully error-corrected quantum computers, NISQ devices used as coprocessors might have advantages over classical computers for solving some problems of practical interest. For example, variational algorithms using NISQ hardware have potential applications to chemistry, materials science, and optimization [2–10].

In a typical NISQ variational algorithm, we need to estimate expectation values for a specified set of operators $\{O_1, O_2, \dots, O_L\}$ in a quantum state ρ that can be prepared repeatedly using a programmable quantum system. To obtain precise estimates, each operator must be measured many times, and finding a reasonably efficient procedure for extracting the desired information is not easy in general. In this Letter, we consider the special case where each O_j is a Pauli operator; this case is of particular interest for near-term applications.

Suppose we have quantum hardware that produces multiple copies of the n -qubit state ρ . Furthermore, for every copy, we can measure all the qubits independently, choosing at our discretion to measure each qubit in the X , Y , or Z basis. We are given a list of L n -qubit Pauli operators (each one a tensor product of n Pauli matrices), and our task is to estimate the expectation values of all L operators in the state ρ , with an additive error no larger than ϵ for each operator. We would like to perform this task using as few copies of ρ as possible.

If all L Pauli operators have relatively low weight (act nontrivially on only a few qubits), there is a simple

randomized protocol that achieves our goal quite efficiently: For each of M copies of ρ , and for each of the n qubits, we chose uniformly at random to measure X , Y , or Z . Then we can achieve the desired prediction accuracy with high success probability if $M = O(3^w \log L/\epsilon^2)$, assuming that all L operators on our list have weight no larger than w [11,12]. If the list contains high-weight operators, however, this randomized method is not likely to succeed unless M is very large.

In this Letter, we describe a deterministic protocol for estimating Pauli-operator expectation values that always performs at least as well as the randomized protocol and performs much better in some cases. This deterministic protocol is constructed by *derandomizing* the randomized protocol. The key observation is that we can compute a lower bound on the probability that randomized measurements on M copies successfully achieve the desired error ϵ for every one of our L target Pauli operators. Furthermore, we can compute this lower bound even when the measurement protocol is partially deterministic and partially randomized; that is, when some of the measured single-qubit Pauli operators are fixed, and others are still sampled uniformly from $\{X, Y, Z\}$.

Hence, starting with the fully randomized protocol, we can proceed step by step to replace each randomized single-qubit measurement by a deterministic one, taking care in each step to ensure that the new partially randomized protocol, with one additional fixed measurement, has success probability at least as high as the preceding protocol. When all measurements have been fixed, we have a fully

deterministic protocol. In numerical experiments, we find that this deterministic protocol substantially outperforms randomized protocols [12–16]. The improvement is especially significant when the list of target observables includes operators with relatively high weight. Further performance gains are possible by executing (at least) linear-depth circuits before measurements [17–20]. Such procedures do, however, require deep quantum circuits. In contrast, our protocol only requires single-qubit Pauli measurements, which are more amenable to execution on near-term devices.

The manuscript is organized as follows. We first provide some statistical background, explain the randomized measurement protocol, then analyze the derandomization procedure. We then provide numerical results showing how the derandomized protocol improves on previous methods. We conclude with remarks and outlooks. Further examples and details of proofs are in the Supplemental Material [21].

Statistical background.—Let ρ be a fixed, but unknown, quantum state on n qubits. We want to accurately predict L expectation values

$$\omega_\ell(\rho) = \text{tr}(O_{\mathbf{o}_\ell}\rho) \quad \text{for } 1 \leq \ell \leq L, \quad (1)$$

where each $O_{\mathbf{o}_\ell} = \sigma_{\mathbf{o}_\ell[1]} \otimes \cdots \otimes \sigma_{\mathbf{o}_\ell[n]}$ is a tensor product of single-qubit Pauli matrices, i.e., $\mathbf{o}_\ell = [\mathbf{o}_\ell[1], \dots, \mathbf{o}_\ell[n]]$ with $\mathbf{o}_\ell[k] \in \{I, X, Y, Z\}$. To extract meaningful information, we perform M (single-shot) Pauli measurements on independent copies of ρ . There are 3^n possible measurement choices. Each of them is characterized by a full-weight Pauli string $\mathbf{p}_m \in \{X, Y, Z\}^n$ and produces a random string of n outcome signs $\mathbf{q}_m \in \{\pm 1\}^n$.

Not every Pauli measurement \mathbf{p}_m ($1 \leq m \leq M$) provides actionable advice about every target observable \mathbf{o}_ℓ ($1 \leq \ell \leq L$). The two must be compatible in the sense that the latter corresponds to a marginal of the former; i.e., it is possible to obtain \mathbf{o}_ℓ from \mathbf{p}_m by replacing some local nonidentity Pauli matrices with I . If this is the case, we write $\mathbf{o}_\ell \triangleright \mathbf{p}_m$ and say that measurement \mathbf{p}_m “hits” target observable \mathbf{o}_ℓ . For instance, $[X, I], [I, X], [X, X] \triangleright [X, X]$, but $[Z, I], [I, Z], [Z, Z] \not\triangleright [X, X]$. We can approximate each $\omega_\ell(\rho)$ by empirically averaging (appropriately marginalized) measurement outcomes that belong to Pauli measurements that hit \mathbf{o}_ℓ ,

$$\hat{\omega}_\ell = \frac{1}{h(\mathbf{o}_\ell; [\mathbf{p}_1, \dots, \mathbf{p}_M])} \sum_{m: \mathbf{o}_\ell \triangleright \mathbf{p}_m} \prod_{j: \mathbf{o}_\ell[j] \neq I} \mathbf{q}_m[j], \quad (2)$$

where $h(\mathbf{o}_\ell; [\mathbf{p}_1, \dots, \mathbf{p}_M]) = \sum_{m=1}^M \mathbf{1}\{\mathbf{o}_\ell \triangleright \mathbf{p}_m\} \in \{0, 1, \dots, M\}$ counts how many Pauli measurements hit target observable \mathbf{o}_ℓ .

It is easy to check that each $\hat{\omega}_\ell$ exactly reproduces $\omega_\ell(\rho)$ in expectation [provided that $h(\mathbf{o}_\ell; \mathbf{P}) \geq 1$]. Moreover, the probability of a large deviation improves exponentially with the number of hits.

Lemma 1. (Confidence bound). Fix $\varepsilon \in (0, 1)$ (accuracy) and $1 - \delta \in (0, 1)$ (confidence). Suppose that Pauli observables $\mathbf{O} = [\mathbf{o}_1, \dots, \mathbf{o}_L]$ and Pauli measurements $\mathbf{P} = [\mathbf{p}_1, \dots, \mathbf{p}_M]$ are such that

$$\text{Conf}_\varepsilon(\mathbf{O}; \mathbf{P}) := \sum_{\ell=1}^L \exp\left(-\frac{\varepsilon^2}{2} h(\mathbf{o}_\ell; \mathbf{P})\right) \leq \frac{\delta}{2}. \quad (3)$$

Then, the associated empirical averages (2) obey

$$|\hat{\omega}_\ell - \omega_\ell(\rho)| \leq \varepsilon \quad \text{for all } 1 \leq \ell \leq L \quad (4)$$

with probability (at least) $1 - \delta$.

See Supplemental Material Sec. B.1 for a detailed derivation [21]. We call the function defined in Eq. (3) the “confidence bound.” It is a statistically sound summary parameter that checks whether a set of Pauli measurements (\mathbf{P}) allows for confidently predicting a collection of Pauli observables (\mathbf{O}) up to accuracy ε each.

Randomized Pauli measurements.—Intuitively speaking, a small confidence bound (3) implies a good Pauli estimation protocol. But how should we choose our M Pauli measurements (\mathbf{P}) in order to achieve $\text{Conf}_\varepsilon(\mathbf{O}; \mathbf{P}) \leq \delta/2$? The randomized measurement toolbox [12, 13, 16, 22, 23] provides a perhaps surprising answer to this question. Let $w(\mathbf{o}_\ell)$ denote the *weight* of Pauli observable \mathbf{o}_ℓ , i.e., the number of qubits on which the observable acts nontrivially: $w(\mathbf{o}_\ell) = \sum_{k=1}^n \mathbf{1}\{\mathbf{o}_\ell[k] \neq I\}$. These weights capture the probability of hitting \mathbf{o}_ℓ with a completely random measurement string: $\text{Prob}_{\mathbf{p}}[\mathbf{o}_\ell \triangleright \mathbf{p}] = 1/3^{w(\mathbf{o}_\ell)}$. In turn, a total of M randomly selected Pauli measurements will, on average, achieve $\mathbb{E}_{\mathbf{p}}[h(\mathbf{o}_\ell; \mathbf{P})] = M/3^{w(\mathbf{o}_\ell)}$ hits, regardless of the actual Pauli observable \mathbf{o}_ℓ in question. This insight allows us to compute expectation values of the confidence bound (3)

$$\mathbb{E}_{\mathbf{p}}[\text{Conf}_\varepsilon(\mathbf{O}; \mathbf{P})] = \sum_{\ell=1}^L (1 - \nu/3^{w(\mathbf{o}_\ell)})^M, \quad (5)$$

where $\nu = 1 - \exp(-\varepsilon^2/2) \in (0, 1)$. Each of the L terms is exponentially suppressed in $\varepsilon^2 M/3^{w(\mathbf{o}_\ell)}$. Concrete realizations of a randomized measurement protocol are extremely unlikely to deviate substantially from this expected behavior (see, e.g., [11]). Combined with Lemma 1, this observation implies a powerful error bound.

Theorem 1. (Theorem 3 in Ref. [11].)—Empirical averages (2) obtained from M randomized Pauli measurements allow for ε -accurately predicting L Pauli expectation values $\text{tr}(O_{\mathbf{o}_1}\rho), \dots, \text{tr}(O_{\mathbf{o}_L}\rho)$ up to additive error ε given that $M \propto \log(L) \max_\ell 3^{w(\mathbf{o}_\ell)}/\varepsilon^2$.

In particular, order $\log(L)$ randomized Pauli measurements suffice for estimating any collection of L low-weight Pauli observables. It is instructive to compare this result

to other powerful statements about randomized measurements, most notably the ‘‘classical shadow’’ paradigm [12,16]. For Pauli observables and Pauli measurements, the two approaches are closely related. The estimators (2) are actually simplified variants of the classical shadow protocol (in particular, they do not require median of means prediction) and the requirements on M are also comparable. This is no coincidence; information-theoretic lower bounds from [12] assert that there are scenarios where the scaling $M \propto \log(L) \max_{\ell} 3^{w(\mathbf{o}_{\ell})} / \varepsilon^2$ is asymptotically optimal and cannot be avoided.

Nevertheless, this does not mean that randomized measurements are *always* a good idea. High-weight observables do pose an immediate challenge, because it is extremely unlikely to hit them by chance alone.

Derandomized Pauli measurements.—The main result of this Letter is a procedure for identifying ‘‘good’’ Pauli measurements that allow for accurately predicting many (fixed) Pauli expectation values. This procedure is designed to interpolate between two extremes: (i) completely randomized measurements (good for predicting many local observables) and (ii) completely deterministic measurements that directly measure observables sequentially (good for predicting few global observables).

Note that we can efficiently compute concrete confidence bounds (3), as well as expected confidence bounds averaged over all possible Pauli measurements (5). Combined, these two formulas also allow us to efficiently compute expected confidence bounds for a list of measurements that is partially deterministic and partially randomized. Suppose that \mathbf{P}^{\sharp} subsumes deterministic assignments for the first $(m-1)$ Pauli measurements, as well as concrete choices for the first $(k-1)$ Pauli labels of the m th measurement, see Fig. 1 (center). There are three possible choices for the next Pauli assignment: $\mathbf{P}^{\sharp}[k, m] = W$ with $W = X, Y, Z$. For each choice, we can explicitly compute the resulting conditional expectation value,

$$\begin{aligned} & \mathbb{E}_{\mathbf{P}}[\text{Conf}_{\varepsilon}(\mathbf{O}; \mathbf{P}) | \mathbf{P}^{\sharp}, \mathbf{P}[k, m] = W] \\ &= \sum_{\ell=1}^L \exp\left(-\frac{\varepsilon^2}{2} \sum_{m'=1}^{m-1} \prod_{k'=1}^n \mathbf{1}\{\mathbf{o}_{\ell}[k'] \triangleright \mathbf{P}^{\sharp}[k', m']\}\right) \\ & \quad \times \left(1 - \nu \frac{\mathbf{1}\{\mathbf{o}_{\ell}[k] \triangleright W\}}{3^{w_{\mathbf{k}}(\mathbf{o}_{\ell})}} \prod_{k'=1}^{k-1} \mathbf{1}\{\mathbf{o}_{\ell}[k'] \triangleright \mathbf{P}^{\sharp}[k', m]\}\right) \\ & \quad \times (1 - \nu 3^{-w(\mathbf{o}_{\ell})})^{M-m}, \end{aligned} \quad (6)$$

where $\nu = 1 - \exp(-\varepsilon^2/2)$, $w_{\mathbf{k}}(\mathbf{o}_{\ell}) = w(\mathbf{o}_{\ell}[k+1], \dots, \mathbf{o}_{\ell}[n])$ and $\mathbf{o}_{\ell}[k'] \triangleright \mathbf{P}^{\sharp}[k', m]$ if $\mathbf{o}_{\ell}[k'] = I$ or $\mathbf{o}_{\ell}[k'] = \mathbf{P}^{\sharp}[k', m]$. This formula allows us to build deterministic measurements one Pauli label at a time.

We start by envisioning a collection of M completely random n -qubit Pauli measurements. That is, each Pauli label is random and Eq. (5) captures the expected confidence bound averaged over *all* 3^{nM} assignments. There are three possible choices for the first label in the first Pauli measurement: $\mathbf{P}[1, 1] = X$, $\mathbf{P}[1, 1] = Y$, and $\mathbf{P}[1, 1] = Z$. At least one concrete choice does not further increase the confidence bound averaged over all remaining Pauli signs,

$$\begin{aligned} & \min_{W \in \{X, Y, Z\}} \mathbb{E}_{\mathbf{P}}[\text{Conf}_{\varepsilon}(\mathbf{O}; \mathbf{P}) | \mathbf{P}[1, 1] = W] \\ & \leq \frac{1}{3} \sum_{W \in \{X, Y, Z\}} \mathbb{E}_{\mathbf{P}}[\text{Conf}_{\varepsilon}(\mathbf{O}; \mathbf{P}) | \mathbf{P}[1, 1] = W] \\ & = \mathbb{E}_{\mathbf{P}}[\text{Conf}_{\varepsilon}(\mathbf{O}; \mathbf{P})]. \end{aligned} \quad (7)$$

Crucially, Eq. (6) allows us to efficiently identify a minimizing assignment

$$\mathbf{P}^{\sharp}[1, 1] = \underset{W \in \{X, Y, Z\}}{\text{argmin}} \mathbb{E}_{\mathbf{P}}[\text{Conf}_{\varepsilon}(\mathbf{O}; \mathbf{P}) | \mathbf{P}[1, 1] = W]. \quad (8)$$

Doing so replaces an initially random single-qubit measurement setting by a concrete Pauli label that minimizes the conditional expectation value over all remaining (random)

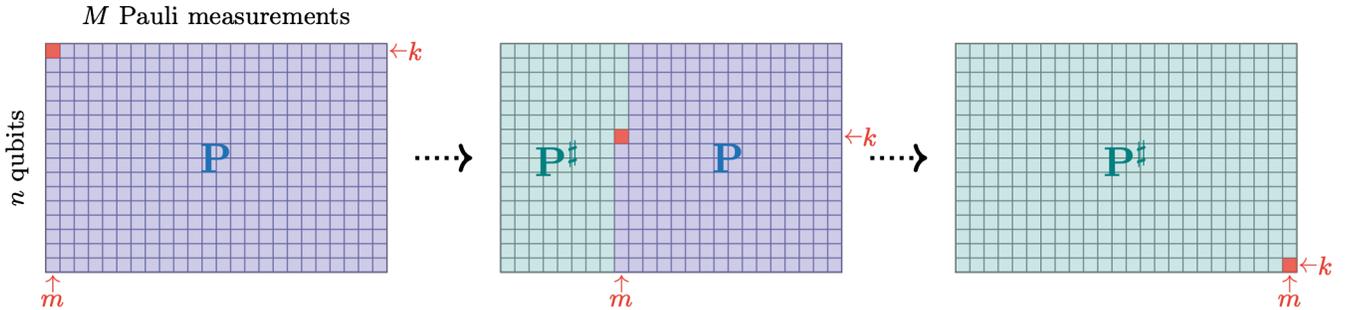


FIG. 1. Illustration of the derandomization algorithm (Algorithm 1): We envision M randomized n -qubit measurements as a two-dimensional array composed of $n \times M$ Pauli labels. Blue squares are place holders for random Pauli labels, while green squares denote deterministic assignments (either X , Y , or Z). Starting with a completely unspecified array (left), the algorithm iteratively checks how a concrete Pauli assignment (red square) affects the confidence bound [Eq. (3)] averaged over all remaining assignments. A simple update rule [Eq. (8)] replaces the initially random label with a deterministic assignment that keeps the remaining confidence bound expectation as small as possible (center). Once the entire grid is traversed, no randomness is left (right) and the algorithm outputs M deterministic n -qubit Pauli measurements.

Algorithm 1. The derandomization algorithm proposed in this work for finding an efficient scheme for measuring a collection of n -qubit Pauli observables.

Derandomization.

Input: measurement budget M , accuracy ε , and L n -qubit Pauli observables $\mathbf{O} = [\mathbf{o}_1, \dots, \mathbf{o}_L]$.

Output: M Pauli measurements $\mathbf{P}^\# \in \{X, Y, Z\}^{n \times M}$.

```

1  function DERANDOMIZATION ( $\mathbf{O}, M, \varepsilon$ )
2      initialize  $\mathbf{P}^\# = [[]]$  (empty  $n \times M$  array)
3      for  $m = 1$  to  $M$  do  $\triangleright$  loop over measurements
4          for  $k = 1$  to  $n$  do  $\triangleright$  loop over qubits
5              for  $W = X, Y, Z$  do compute
6                   $f(W) = \mathbb{E}_{\mathbf{P}}[\text{Conf}_\varepsilon(\mathbf{O}; \mathbf{P})]$ 
6                       $\mathbf{P}^\#, \mathbf{P}[k, m] = W$ 
7                  [see Eq. (6) for a precise formula]
8                   $\mathbf{P}^\#[k, m] \leftarrow \text{argmin}_{W \in \{X, Y, Z\}} f(W)$ 
9      output  $\mathbf{P}^\# \in \{X, Y, Z\}^{n \times M}$ 
    
```

assignments. This procedure is known as derandomization [24–26] and can be iterated. Figure 1 provides visual guidance, while pseudo-code can be found in Algorithm 1. There are a total of $n \times M$ iterations. Step (k, m) is contingent on comparing three conditional expectation values $\mathbb{E}_{\mathbf{P}}[\text{Conf}_\varepsilon(\mathbf{O}; \mathbf{P}) | \mathbf{P}^\#, \mathbf{P}[k, m] = W]$ and assigning the Pauli label that achieves the smallest score. These update rules are constructed to ensure that (appropriate modifications of) Eq. (7) remain valid throughout the procedure. Combining all of them implies the following rigorous statement about the resulting Pauli measurements $\mathbf{P}^\#$.

Theorem 2. (Derandomization promise).—Algorithm 1 is guaranteed to output Pauli measurements $\mathbf{P}^\#$ with below average confidence bound: $\text{Conf}_\varepsilon(\mathbf{O}; \mathbf{P}^\#) \leq \mathbb{E}_{\mathbf{P}}[\text{Conf}_\varepsilon(\mathbf{O}; \mathbf{P})]$.

We see that derandomization produces deterministic Pauli measurements that perform at least as favorably as (averages of) randomized measurement protocols. But the actual difference between randomized and derandomized Pauli measurements can be much more pronounced. In the examples we considered, derandomization reduces the measurement budget M by at least an order of magnitude, compared to randomized measurements. Furthermore, because Algorithm 1 implements a greedy update procedure, we have no assurance that our derandomized measurement procedure is globally optimal or even close to optimal. Using dynamic programming, the derandomization algorithm runs in time $\mathcal{O}(nML)$; see Supplemental Material Sec. C 3 for a detailed implementation [21].

Numerical experiments.—The ability to accurately estimate many Pauli observables is an essential subroutine for variational quantum eigensolvers (VQEs) [4,8–10,27]. Randomized Pauli measurements [11,12]—also known as classical shadows in this context—offer a conceptually simple solution that is efficient both in terms of quantum hardware and measurement budget.

Derandomization can and should be viewed as a refinement of the original classical shadows idea. Supported by rigorous theory (Theorem 2), this refinement is only contingent on an efficient classical preprocessing step, namely, running Algorithm 1. It does not incur any extra cost in terms of quantum hardware and classical postprocessing, but can lead to substantial performance gains. Numerical experiments visualized in Ref. [12], Fig. 5, have revealed unconditional improvements of about one order of magnitude for a particular VQE experiment [28] (simulating quantum field theories).

In this section, we present additional numerical studies that support this favorable picture. These address a slight variation of Algorithm 1 that does not require fixing the total measurement budget M in advance. We focus on the “electronic structure problem”: determine the ground-state energy for molecules with unknown electronic structure. This is one of the most promising VQE applications in quantum chemistry and material science. Different encoding schemes—most notably Jordan-Wigner (JW) [29], Bravyi-Kitaev (BK) [30] and parity (P) [30,31]—allow for mapping molecular Hamiltonians to qubit Hamiltonians that correspond to sums of Pauli observables. Several benchmark molecules have been identified whose encoded Hamiltonians are just simple enough for an explicit classical minimization, so that we can compare Pauli estimation techniques with the exact answer.

Figure 2 illustrates one such comparison. We fix a benchmark molecule BeH_2 , a BK encoding and plot the ground-state energy approximation error against the

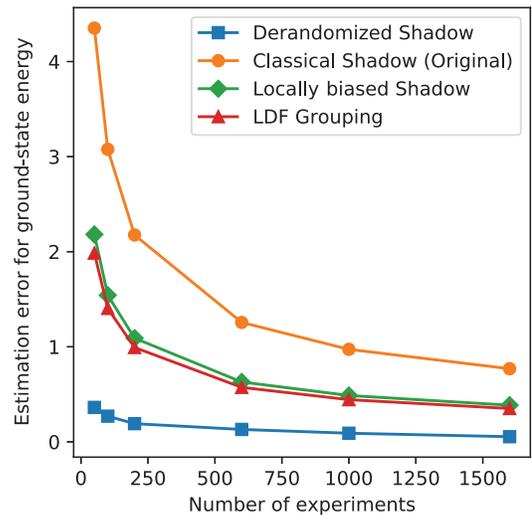


FIG. 2. BeH_2 ground-state energy estimation error (in Hartree) under Bravyi-Kitaev encoding [30] for different measurement schemes: The error for derandomized shadow is the root-mean-squared error (RMSE) over ten independent runs. The error for the other methods shows the RMSE over infinitely many runs and can be evaluated efficiently using the variance of one experiment [14].

TABLE I. Average estimation error using 1000 measurements for different molecules, encodings, and measurement schemes: The first column shows the molecule and the corresponding ground-state electronic (Elec.) energy (in Hartree). We consider the following abbreviations: Derandomized Pauli measurements (Derand.), locally biased classical shadow (Local S.), largest degree first (LDF) heuristic, and original classical shadow (Shadow) [12]

Molecule (E_{GS})	Encodings	Derand.	Local S.	LDF	Shadow
H_2 (-1.86)	<i>JW</i>	0.06	0.13	0.15	0.41
	<i>P</i>	0.03	0.14	0.19	0.48
	<i>BK</i>	0.06	0.14	0.19	0.75
Li H (-8.91)	<i>JW</i>	0.03	0.12	0.23	0.52
	<i>P</i>	0.03	0.16	0.29	0.87
	<i>BK</i>	0.04	0.26	0.27	0.40
BeH_2 (-19.04)	<i>JW</i>	0.06	0.26	0.37	1.29
	<i>P</i>	0.09	0.36	0.49	1.77
	<i>BK</i>	0.06	0.49	0.44	0.97
H_2O (-83.60)	<i>JW</i>	0.12	0.51	1.02	1.68
	<i>P</i>	0.22	0.65	1.63	2.52
	<i>BK</i>	0.20	1.17	1.45	3.25
NH_3 (-66.88)	<i>JW</i>	0.18	0.59	0.94	3.79
	<i>P</i>	0.21	0.83	1.61	2.13
	<i>BK</i>	0.12	0.73	1.45	1.89

number of Pauli measurements. The plot highlights that derandomization outperforms the original classical shadows procedure (randomized Pauli measurements) [12], locally biased classical shadows [12], and another popular technique known as LDF grouping [14,32]. The discrepancy between randomized and derandomized Pauli measurements is particularly pronounced.

This favorable picture extends to a variety of other benchmark molecules and other encoding schemes, see Table I. For a fixed measurement budget, derandomization consistently leads to a smaller estimation error than other state-of-the-art techniques. One could also repeat the measurement scheme found by the derandomization algorithm multiple times to improve the estimation error; see Supplemental Material Sec. C.4 [21]. Finally, we note that in the presence of measurement noise, the various approaches we have considered are likely to suffer about equally, as they were all based on single-qubit Pauli measurements. One could mitigate such noise by incorporating recently proposed noise inversion techniques [33,34].

Conclusion and outlook.—We consider the problem of predicting many Pauli expectation values from few Pauli measurements. Derandomization [24–26] provides an efficient procedure that replaces originally randomized single-qubit Pauli measurements by specific Pauli assignments. The resulting Pauli measurements are deterministic, but inherit *all* advantages of a fully randomized measurement protocol. Furthermore, the derandomization procedure

can accurately capture the fine-grained structure of the observables in question. Predicting molecular ground-state energies based on derandomized Pauli measurements scales favorably and improves upon many existing techniques [11,14,16,32]. Source code for an implementation of the proposed procedure is available at [35].

Randomized measurements have also been used to estimate entanglement entropy [12,36–38], topological invariants [39,40], benchmark physical devices [12,22,41,42], and predict outcomes of physical experiments [43]. Derandomization provides a principled approach for adapting randomized measurement procedures to fine-grained structure and is closely related to an algorithmic technique—multiplicative weight update [44]—commonly used in machine learning and game theory. So far, we have only considered estimations of Pauli observables, but measurement design via derandomization should apply more broadly; we look forward to applying derandomization to other tasks such as estimating non-Pauli observables and entanglement entropies. Additional improvements in performance might be achieved by modifying the cost function $f(W)$ used in Algorithm 1, for example, by greedily assigning more than one single-qubit Pauli measurement in each iteration.

The authors thank Andreas Elben, Stefan Hillmich, Steven T. Flammia, Jarrod McClean, and Lorenzo Pastori for valuable input and inspiring discussions. H.H. is supported by the J. Yang and Family Foundation. J.P. acknowledges funding from the U.S. Department of Energy Office of Science, Office of Advanced Scientific Computing Research, (DE-NA0003525, DE-SC0020290), and the National Science Foundation (PHY-1733907). The Institute for Quantum Information and Matter is a NSF Physics Frontiers Center.

*hsinyuan@caltech.edu

- [1] J. Preskill, Quantum computing in the NISQ era and beyond, *Quantum* **2**, 79 (2018).
- [2] K. Bharti, A. Cervera-Lierta, T.H. Kyaw, T. Haug, S. Alperin-Lea, A. Anand, M. Degroote, H. Heimonen, J.S. Kottmann, T. Menke, W.-K. Mok, S. Sim, L.-C. Kwek, and A. Aspuru-Guzik, Noisy intermediate-scale quantum (NISQ) algorithms, [arXiv:2101.08448](https://arxiv.org/abs/2101.08448).
- [3] M. Cerezo, A. Arrasmith, R. Babbush, S. C. Benjamin, S. Endo, K. Fujii, J.R. McClean, K. Mitarai, X. Yuan, L. Cincio *et al.*, Variational quantum algorithms, [arXiv:2012.09265](https://arxiv.org/abs/2012.09265).
- [4] C. Hempel, C. Maier, J. Romero, J. McClean, T. Monz, H. Shen, P. Jurcevic, B. P. Lanyon, P. Love, R. Babbush, A. Aspuru-Guzik, R. Blatt, and C. F. Roos, Quantum Chemistry Calculations on a Trapped-Ion Quantum Simulator, *Phys. Rev. X* **8**, 031022 (2018).
- [5] H.-Y. Huang, K. Bharti, and P. Rebentrost, Near-term quantum algorithms for linear systems of equations, [arXiv:1909.07344](https://arxiv.org/abs/1909.07344).

- [6] H.-Y. Huang, M. Broughton, M. Mohseni, R. Babbush, S. Boixo, H. Neven, and J.R. McClean, Power of data in quantum machine learning, *Nat. Commun.* **12**, 2631 (2021).
- [7] A. Kandala, A. Mezzacapo, K. Temme, M. Takita, M. Brink, J. M. Chow, and J. M. Gambetta, Hardware-efficient variational quantum eigensolver for small molecules and quantum magnets, *Nature (London)* **549**, 242 (2017).
- [8] P. J. J. O'Malley, R. Babbush, I. D. Kivlichan, J. Romero, J. R. McClean, R. Barends, J. Kelly, P. Roushan, A. Tranter, N. Ding *et al.*, Scalable Quantum Simulation of Molecular Energies, *Phys. Rev. X* **6**, 031007 (2016).
- [9] A. Peruzzo, J. McClean, P. Shadbolt, M.-H. Yung, X.-Q. Zhou, P. J. Love, A. Aspuru-Guzik, and J. L. O'Brien, A variational eigenvalue solver on a photonic quantum processor, *Nat. Commun.* **5**, 4213 (2014).
- [10] G. A. Quantum *et al.*, Hartree-Fock on a superconducting qubit quantum computer, *Science* **369**, 1084 (2020).
- [11] T. J. Evans, R. Harper, and S. T. Flammia, Scalable Bayesian Hamiltonian learning, [arXiv:1912.07636](https://arxiv.org/abs/1912.07636).
- [12] H.-Y. Huang, R. Kueng, and J. Preskill, Predicting many properties of a quantum system from very few measurements, *Nat. Phys.* **16**, 1050 (2020).
- [13] A. Elben, B. Vermersch, C. F. Roos, and P. Zoller, Statistical correlations between locally randomized measurements: A toolbox for probing entanglement in many-body quantum states, *Phys. Rev. A* **99**, 052323 (2019).
- [14] C. Hadfield, S. Bravyi, R. Raymond, and A. Mezzacapo, Measurements of quantum Hamiltonians with locally-biased classical shadows, [arXiv:2006.15788](https://arxiv.org/abs/2006.15788).
- [15] M. Ohliger, V. Nesme, and J. Eisert, Efficient and feasible state tomography of quantum many-body systems, *New J. Phys.* **15**, 015024 (2013).
- [16] M. Pains and A. Kalev, An approximate description of quantum states, [arXiv:1910.10543](https://arxiv.org/abs/1910.10543).
- [17] O. Crawford, B. van Straaten, D. Wang, T. Parks, E. Campbell, and S. Brierley, Efficient quantum measurement of Pauli operators in the presence of finite sampling error, *Quantum* **5**, 385 (2021).
- [18] W. J. Huggins, J. R. McClean, N. C. Rubin, Z. Jiang, N. Wiebe, K. B. Whaley, and R. Babbush, Efficient and noise resilient measurements for quantum chemistry on near-term quantum computers, *npj Quantum Inf.* **7**, 23 (2021).
- [19] A. F. Izmaylov, T.-C. Yen, R. A. Lang, and V. Verteletskyi, Unitary partitioning approach to the measurement problem in the variational quantum eigensolver method, *J. Chem. Theory Comput.* **16**, 190 (2020).
- [20] T.-C. Yen and A. F. Izmaylov, Cartan sub-algebra approach to efficient measurements of quantum observables, [arXiv:2007.01234](https://arxiv.org/abs/2007.01234).
- [21] See Supplemental Material at <http://link.aps.org/supplemental/10.1103/PhysRevLett.127.030503> for additional proofs of the theorems and details of the numerical experiments.
- [22] A. Elben, R. Kueng, H.-Y. R. Huang, R. van Bijnen, C. Kokail, M. Dalmonte, P. Calabrese, B. Kraus, J. Preskill, P. Zoller, and B. Vermersch, Mixed-State Entanglement from Local Randomized Measurements, *Phys. Rev. Lett.* **125**, 200501 (2020).
- [23] M. Ohliger, V. Nesme, and J. Eisert, Efficient and feasible state tomography of quantum many-body systems, *New J. Phys.* **15**, 015024 (2013).
- [24] N. Alon and J. H. Spencer, *The Probabilistic Method*, 3rd ed. Wiley-Interscience Series in Discrete Mathematics and Optimization (Wiley, New York, 2008).
- [25] R. Motwani and P. Raghavan, *Randomized Algorithms* (Cambridge University Press, Cambridge, England, 1995).
- [26] V. V. Vazirani, *Approximation Algorithms* (Springer, New York, 2001).
- [27] A. Kandala, A. Mezzacapo, K. Temme, M. Takita, M. Brink, J. M. Chow, and J. M. Gambetta, Hardware-efficient variational quantum eigensolver for small molecules and quantum magnets, *Nature (London)* **549**, 242 (2017).
- [28] C. Kokail, C. Maier, R. van Bijnen, T. Brydges, M. K. Joshi, P. Jurcevic, C. A. Muschik, P. Silvi, R. Blatt, C. F. Roos *et al.*, Self-verifying variational quantum simulation of lattice models, *Nature (London)* **569**, 355 (2019).
- [29] P. Jordan and E. Wigner, Über das paulische äquivalenzverbot, *Z. Phys.* **47**, 631 (1928).
- [30] S. B. Bravyi and A. Y. Kitaev, Fermionic quantum computation, *Ann. Phys. (Amsterdam)* **298**, 210 (2002).
- [31] J. T. Seeley, M. J. Richard, and P. J. Love, The Bravyi-Kitaev transformation for quantum computation of electronic structure, *J. Chem. Phys.* **137**, 224109 (2012).
- [32] V. Verteletskyi, T.-C. Yen, and A. F. Izmaylov, Measurement optimization in the variational quantum eigensolver using a minimum clique cover, *J. Chem. Phys.* **152**, 124114 (2020).
- [33] S. Chen, W. Yu, P. Zeng, and S. T. Flammia, Robust shadow estimation, [arXiv:2011.09636](https://arxiv.org/abs/2011.09636).
- [34] D. E. Koh and S. Grewal, Classical shadows with noise, [arXiv:2011.11580](https://arxiv.org/abs/2011.11580).
- [35] H.-Y. Huang, R. Kueng, and J. Preskill, Source code for the derandomization procedure, <https://github.com/momohuang/predicting-quantum-properties> (accessed 2021).
- [36] T. Brydges, A. Elben, P. Jurcevic, B. Vermersch, C. Maier, B. P. Lanyon, P. Zoller, R. Blatt, and C. F. Roos, Probing rényi entanglement entropy via randomized measurements, *Science* **364**, 260 (2019).
- [37] A. Rath, R. van Bijnen, A. Elben, P. Zoller, and B. Vermersch, Importance sampling of randomized measurements for probing entanglement, [arXiv:2102.13524](https://arxiv.org/abs/2102.13524).
- [38] V. Vitale, A. Elben, R. Kueng, A. Neven, J. Carrasco, B. Kraus, P. Zoller, P. Calabrese, B. Vermersch, and M. Dalmonte, Symmetry-resolved dynamical purification in synthetic quantum matter, [arXiv:2101.07814](https://arxiv.org/abs/2101.07814).
- [39] Z.-P. Cian, H. Dehghani, A. Elben, B. Vermersch, G. Zhu, M. Barkeshli, P. Zoller, and M. Hafezi, Many-Body Chern Number from Statistical Correlations of Randomized Measurements, *Phys. Rev. Lett.* **126**, 050501 (2021).
- [40] 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.* **6**, eaaz3666 (2020).
- [41] J. Choi, A. L. Shaw, I. S. Madjarov, X. Xie, J. P. Covey, J. S. Cotler, D. K. Mark, H.-Y. Huang, A. Kale, H. Pichler, F. G. S. L. Brandão, S. Choi, and M. Endres, Emergent

- randomness and benchmarking from many-body quantum chaos, [arXiv:2103.03535](#).
- [42] E. Knill, D. Leibfried, R. Reichle, J. Britton, R. B. Blakestad, J. D. Jost, C. Langer, R. Ozeri, S. Seidelin, and D. J. Wineland, Randomized benchmarking of quantum gates, *Phys. Rev. A* **77**, 012307 (2008).
- [43] H.-Y. Huang, R. Kueng, and J. Preskill, Information-theoretic bounds on quantum advantage in machine learning, *Phys. Rev. Lett.* **126**, 190505 (2021).
- [44] S. Arora, E. Hazan, and S. Kale, The multiplicative weights update method: A meta-algorithm and applications, *Theory Comput.* **8**, 121 (2012).