Improving Classical Shadows with Grouping Strategies

Author: Marçal Herraiz–Bayó.
Facultat de Física, Universitat de Barcelona, Diagonal 645, 08028 Barcelona, Spain.[*]

Advisors: Berta Casas–Font[1], Sergi Masot–Llima[1]. Tutor: Bruno Juliá–Díaz[2].[†]
(Dated: June 13, 2025)

Abstract: Efficiently extracting information from quantum systems is a key challenge in quantum computing. This thesis explores the combination of two complementary techniques—*classical shadows* and *grouping*—to improve quantum measurement strategies under resource constraints. Classical shadows enable the prediction of many properties of a quantum state from a small number of measurements, while grouping strategies reduce the number of measurements needed by exploiting commutativity among observables. We implement a hybrid method, *Shadow–Grouping*, that unifies both approaches to enhance measurement efficiency on current quantum devices. We also demonstrate its effectiveness by estimating the ground state energies of H_2 and LiH molecules: compared to standard classical shadows, our results show that Shadow–Grouping achieves up to 18-fold gains in accuracy and reaches chemical precision using orders of magnitude fewer measurements.
Keywords: Quantum computing, classical shadows, grouping, many-body Hamiltonians.
SDGs: 7, 9, 13.

I. INTRODUCTION

Predicting the physical properties of a system based on empirical observations lies at the heart of scientific progress, and quantum physics is no exception. In this discipline, *observables* of a quantum state are of vital importance, as their expectation values encode crucial information about the state, such as its energy.

However, the nature of quantum mechanics is probabilistic—governed by Born's rule—and destructive due to wavefunction collapse. As a consequence, in practical settings, expectation values have to be estimated from many experiments on independent and identically prepared quantum states—a resource-intensive process.

The earliest method for learning the properties of quantum systems was *quantum state tomography*, which attempts to reconstruct the full state from measurement data. Unfortunately, this approach becomes infeasible for large systems, as the number of parameters required to describe an n-qubit state scales exponentially as 2^n.

To circumvent this scalability barrier, *classical shadows* [1, 2] emerged as a powerful alternative, and showed that full-state reconstruction is often unnecessary for concrete tasks, such as estimating expectation values. This method allows us to predict many properties of a quantum state from a small number of measurements, without requiring full-state reconstruction.

A complementary strategy is that of *observable grouping* [3], which clusters commuting observables—such as terms in a Hamiltonian—so that they can be measured simultaneously. As a consequence, this technique reduces the number of distinct measurements required and thus improves measurement efficiency.

These two ideas—classical shadows and grouping—are especially valuable in estimating the energy of quantum many-body Hamiltonians. However, on today's noisy intermediate-scale quantum (NISQ) hardware, state preparation is imperfect and measurements are limited. Thus, performing this task accurately and efficiently poses a major challenge.

In this work, we follow the recent trend of efforts to unify these two approaches under a *Shadow–Grouping* framework [4, 5]. We numerically implement such a hybrid method using the `PennyLane` library [6] and demonstrate its effectiveness by estimating the ground state energies of the H_2 and LiH molecules [7]. Our results showcase a significant improvement in measurement efficiency over previous measurement strategies.

II. THEORETICAL FRAMEWORK

A. Grouping Techniques

A common bottleneck in quantum computing arises when estimating the expectation values of a set of observables $\{O_i\}_{i=1}^M$: typically, each observable must be measured individually. For example, if a many-body Hamiltonian $H = \sum_i H_i$ is expressed as a sum of measurable terms, estimating $\langle H \rangle$ generally requires separate measurements to estimate each term H_i. *Grouping techniques* [3] aim to mitigate this inefficiency by identifying sets of commuting observables that can be measured simultaneously, thus reducing the number of distinct measurements that have to be performed.

As an example, let O_1 and O_2 be two commuting observables. Then, they share a *common eigenbasis* $\{|\psi_i\rangle\}_i$,

[*]Electronic address: mherraba9@alumnes.ub.edu
[†][1]Barcelona Supercomputing Center, Plaça Eusebi Güell, 1-3, 08034 Barcelona, Spain; [2]Departament de Física Quàntica i Astrofísica and Institut de Ciències del Cosmos (ICCUB), Universitat de Barcelona, Martí i Franquès 1, E–08028 Barcelona, Spain.

satisfying $O_k|\psi_i\rangle = \lambda_i^{(k)}|\psi_i\rangle$ for $k = 1, 2$. Given a state $|\Psi\rangle$, the expectation value of each O_k in this state is

$$\langle\Psi|O_k|\Psi\rangle = \sum_i \lambda_i^{(k)} |\langle\psi_i|\Psi\rangle|^2.$$

Therefore, if the eigenvalues $\lambda_i^{(k)}$ are known, performing a single measurement of the probabilities $|\langle\psi_i|\Psi\rangle|^2$ in the common eigenbasis allows us to calculate the expectation value of both observables O_1 and O_2 simultaneously.

In practice, one does not consider full commutativity of the observables, but rather *qubit-wise commutativity*, which is a stronger condition and thus more restrictive. For two tensor product operators $A = A_1 \otimes \cdots \otimes A_n$ and $B = B_1 \otimes \cdots \otimes B_n$ on an n-qubit system, we say that they *qubit-wise commute* (QWC) if $[A_i, B_i] = 0$ for all $i = 1, \ldots, n$. For instance, $X \otimes Y$ and $X \otimes \mathbb{I}$ do QWC, where $\mathbb{I}$ is the identity operator. If two operators QWC, then they commute, but the converse is not true; for example, $X \otimes X$ and $Y \otimes Y$ commute but do not QWC.

Given a set $\{O_i\}_{i=1}^M$ of observables, we aim to partition it into subsets such that all members of each subset pairwise QWC. The fewer the subsets required, the fewer distinct measurements will be needed to estimate all expectation values. However, qubit-wise commutativity is a reflexive and symmetric relation, but not transitive, meaning that such a partition is generally not unique.

To find such a partition, it is useful to reformulate the problem in terms of *graphs*. Each node of the graph will represent an observable of our collection, and we will connect two such nodes if and only if their corresponding observables QWC. In this representation, finding an optimal grouping is reduced to a well-known problem in graph theory known as the *Minimum Clique Cover* (MCC) problem. A *clique* is a subset of nodes in the graph that are all connected to each other. Figure 1 illustrates this idea, where each color represents a clique. The cover (c) uses three cliques, whereas (b) uses two and is an MCC, since no one-clique cover exists (because not all nodes of the graph are connected to each other).

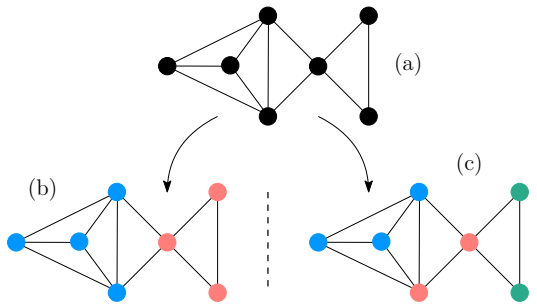

Figure 1: A graph (a) and two clique covers of it ((b), (c)) with a different number of cliques ((b): 2, minimal; (c): 3).

The MCC problem is generally NP-hard [3], and thus we rely on heuristic techniques to find approximate solutions. Accordingly, finding an optimal grouping from a collection of observables is also NP-hard.

B. Classical Shadows

While grouping techniques reduce the total number of measurements needed by combining QWC observables, *classical shadows* [1] aim to estimate many properties of a quantum state using as few measurements as possible, without full tomography or prior grouping being needed.

Given a fixed (but unknown) quantum state ρ of an n-qubit system and a collection $\{O_i\}_{i=1}^M$ of observables, the goal is to estimate the expectation values $o_i(\rho) := \mathrm{Tr}(O_i\rho)$, which are linear functions of the state. Classical shadows provide a framework for efficiently estimating these quantities using randomized measurements and classical post-processing.

To create a classical shadow of a state ρ, fix a collection $\mathcal{U}$ of unitary matrices. Then, randomly select a unitary $U \in \mathcal{U}$ and evolve the state by U, thus obtaining $U\rho U^\dagger$. Performing a measurement in the computational basis yields a n-bit string outcome $b \in \{0,1\}^n$. We can then apply the inverse of U to the resulting computational basis state and store the resulting density matrix $U^\dagger|b\rangle\langle b|U$ in classical memory, which is called a *snapshot* of the state. The expected value of these snapshots over both the randomness of U and the measurement outcomes b defines a *quantum channel* $\mathcal{M}$ that acts on the state ρ:

$$\mathcal{M}(\rho) := \mathop{\mathbb{E}}_{U\in\mathcal{U}} \sum_{b\in\{0,1\}^n} \langle b|U\rho U^\dagger|b\rangle U^\dagger|b\rangle\langle b|U.$$

If the operators in $\mathcal{U}$ form an operator basis of the system's Hilbert space—i.e., the ensemble $\mathcal{U}$ is *tomographically complete*—then this channel $\mathcal{M}$ is invertible. The original state can then be recovered in expectation as

$$\rho = \mathop{\mathbb{E}}_{\text{outcomes } |b\rangle} \left[\mathcal{M}^{-1}\left(U^\dagger|b\rangle\langle b|U\right)\right].$$

Repeating this process N times results in N snapshots of ρ, which together form the *classical shadow* of ρ:

$$S(\rho; N) = \left\{\hat{\rho}_k = \mathcal{M}^{-1}\left(U_k^\dagger|b_k\rangle\langle b_k|U_k\right)\right\}_{k=1,\ldots,N}.$$

With this approximation of the state, we can estimate the expectation values of the observables O_i as $\hat{o}_i = \mathrm{Tr}(O_i\hat{\rho}_k)$, for each snapshot $\hat{\rho}_k$ of the state, and then use empirical mean estimation to approximate o_i.

In this work, we fix $\mathcal{U}$ to be the set of single-qubit Clifford unitaries. These unitaries are defined as the ones that map Pauli matrices to Pauli matrices under conjugation. Thus, applying a single-qubit Clifford unitary U is equivalent to measuring the qubit in the X, Y, or Z Pauli basis, which we will denote by P_U. In this setting, Huang et al. [1] show that the inverted channel has a simple tensor product form:

$$\hat{\rho} = \mathcal{M}^{-1}\left(U^\dagger|b\rangle\langle b|U\right) = \bigotimes_{j=1}^n \left(3U_j^\dagger|b_j\rangle\langle b_j|U_j - \mathbb{I}\right), \quad (1)$$

where $U = U_1 \otimes \cdots \otimes U_n \in \mathcal{U}$ is the randomly chosen single-qubit Clifford gate applied to ρ and $b =$

$(b_1, \ldots, b_n) \in \{0,1\}^n$ is the n-bit string resulting from the computational basis measurement.

Suppose that the observable O we want to estimate is a *Pauli word*, i.e., it has the form $O = P_1 \otimes \cdots \otimes P_n$, where each $P_j \in \{\mathbb{I}, X, Y, Z\}$ is a Pauli operator acting on the qubit j. Using Eq. (1), the expected value of O in the snapshot state $\hat{\rho}$ is given by

$$\mathrm{Tr}(O\hat{\rho}) = \prod_{j=1}^{n} \mathrm{Tr}\left(3P_j U_j^\dagger |b_j\rangle\langle b_j| U_j\right) = \prod_{j:\mathbb{I}\neq P_j = P_{U_j}} 3(1-2b_j), \quad (2)$$

where P_{U_j} denotes the Pauli basis in which qubit j was measured. Therefore, calculating the expectation value of a Pauli word in a snapshot state reduces to counting the number of matches between the Pauli operators in the word and the measurement bases in the shadow, and then multiplying by the appropriate factor of $3(1-2b_j) = \pm 3$.

C. Shadow–Grouping

In the classical shadows framework, we are often interested in estimating the energy of a quantum state ρ with respect to a given Hamiltonian H, i.e., computing $\langle H\rangle = \mathrm{Tr}(H\rho)$. In practice, however, one is constrained by a limited number N of measurements—commonly referred to as the *measurement budget* or the number of *shots*. Furthermore, for many-body Hamiltonians, direct energy measurement is not feasible. Instead, H is expressed as a weighted sum of Pauli words:

$$H = \sum_{i=1}^{M} h_i O^{(i)}, \quad \text{with } O^{(i)} \in \{\mathbb{I}, X, Y, Z\}^{\otimes n} \quad (3)$$

and $h_i \in \mathbb{R}$. Without loss of generality, we assume that $O^{(i)} \neq \mathbb{I}^{\otimes n}$ for all i. The computation of this decomposition of H is beyond the scope of this work, but we remark that it is typically obtained via Hartree–Fock methods followed by fermion-to-qubit mappings, such as Jordan–Wigner (JW) or Bravyi–Kitaev transformations.

A known limitation of the classical shadows method is that for a given Pauli word $O = P_1 \otimes \cdots \otimes P_n$, only the snapshots whose measurement bases match that of O can be used to estimate $\langle O\rangle$ using Eq. (2). All other measurements are effectively discarded, which is inefficient (this was addressed by Huang et al. in [2]).

In this context, grouping the terms $O^{(i)}$ from Eq. (3) into qubit-wise commuting subsets improves the efficiency of classical shadows. However, two challenges arise: (1) finding an optimal QWC-based grouping is NP-hard, and (2) some terms in (3) contribute more strongly to $\langle H\rangle$ and should therefore be prioritized in measurement to obtain an accurate estimation of the energy. This is particularly relevant in quantum chemistry, where high precision in energy estimates is required.

To address these challenges, the *Shadow–Grouping* strategy offers a way to incorporate grouping into the classical shadows framework. The key idea is to determine, for each shot k, a Pauli word $Q^{(k)}$—called a *measurement setting*—that QWC with as many terms $O^{(i)}$ as possible. Then, we measure in the basis defined by $Q^{(k)}$. Because both $Q^{(k)}$ and $O^{(i)}$ are Pauli words, QWC ensures that the basis used in the classical shadows procedure is compatible with $O^{(i)}$, and hence allows for an estimation of $\langle O^{(i)}\rangle$ using Eq. (2) without discarding the shot.

To prioritize important terms in the Hamiltonian, we assign a weight w_i to each $O^{(i)}$ subject to the following conditions:

(i) It should increase with the magnitude of h_i to reflect the importance of the term in the Hamiltonian.

(ii) It should decrease with the number of times $O^{(i)}$ has already been measured, to promote diversity.

These criteria are satisfied by the following weight function proposed by Gresch et al. [4]:

$$w_i := |h_i| \frac{\sqrt{N_i+1} - \sqrt{N_i}}{\sqrt{N_i(N_i+1)}},$$

where N_i is the number of times $O^{(i)}$ has been estimated so far. Initially, all weights are set to $w_i = |h_i|$. Note that after each shot k, N_i either remains unchanged—if the measurement setting $Q^{(k)}$ does not QWC with $O^{(i)}$—or increases by one otherwise.

To construct the measurement setting $Q^{(k)}$, we follow the greedy algorithm proposed in [4]: in the k-th shot, we initialize $Q^{(k)} = \mathbb{I}^{\otimes n}$ and iterate through the list of Pauli terms $O^{(i)}$ in (3) sorted in decreasing order of weight. At each step, if $O^{(i)}$ QWC with the current $Q^{(k)}$, we update $Q^{(k)}$ by replacing its identity components with those of $O^{(i)}$. This process continues until either all observables have been considered or $Q^{(k)}$ contains no more identity components. Each measurement setting $Q^{(k)}$ is guaranteed to be compatible with at least one term, namely the one with the largest weight at the given shot. The final set of measurement settings $\{Q^{(k)}\}_{k=1}^{N}$ defines the *measurement scheme*. The pseudocode for this procedure is provided in Algorithm 1.

This approach presents three main advantages over previous methods. First, the weight function from the algorithm is highly adaptable, and it need only provide a hierarchy for the terms in the Hamiltonian decomposition. Second, since each measurement setting QWC with at least one term of H, no measurements are discarded during the classical shadows estimation. Lastly, unlike traditional grouping strategies [3], this procedure avoids solving the MCC problem in the preprocessing step.

III. CODE IMPLEMENTATIONS

One of the primary goals of this thesis was to implement the algorithms for: computing the classical shadow

of a quantum state, constructing measurement schemes tailored to a given Hamiltonian decomposition (3), and estimating the energy $\langle H \rangle$ using the Shadow–Grouping procedure. All code and data generated in this work are publicly available at [8].

The pseudocode for obtaining the measurement setting at each shot, adapted from [4], is given in Algorithm 1:

Algorithm 1 Measurement setting at shot k.

Require: Hamiltonian in Pauli word representation as in (3)
Require: Previous measurement settings $\{Q^{(1)}, \ldots, Q^{(k-1)}\}$
Require: Weight function w
$Q^{(k)} \leftarrow \mathbb{I}^{\otimes n}$ ▷ Initialize as identity on n qubits
$w_i \leftarrow w(O^{(i)})$ for all i ▷ Compute weights
while $Q^{(k)}$ has identity components **and** $\max_i w_i > 0$ **do**
  $j \leftarrow \arg\max_i w_i$ ▷ Select index of maximum weight
  $R \leftarrow \{i : O_i^{(j)} \neq \mathbb{I} \text{ and } Q_i^{(k)} = \mathbb{I}\}$
  if $O^{(j)}$ and $Q^{(k)}$ QWC **then**
    $Q_i^{(k)} \leftarrow O_i^{(j)}$ for all $i \in R$ ▷ Update meas. setting
  end if
  $w_j \leftarrow 0$ ▷ Mark $O^{(j)}$ as processed
end while
return $\{Q^{(1)}, \ldots, Q^{(k)}\}$

To illustrate the behavior of this procedure, consider the following toy Hamiltonian on three qubits:

$$H = h_1 X_1 Z_2 + h_2 Y_1 Z_3 + h_3 Z_2 Z_3 + h_4 X_1 Y_2 Z_3, \quad (4)$$

with increasing positive coefficients h_i, i.e., $h_i < h_j$ for $i < j$. The evolution of the weights across multiple shots is shown in Figure 2. Since the term $X_1 Y_2 Z_3$ has the largest initial weight, the first measurement setting $Q^{(1)}$ is made to QWC with it. However, doing so exhausts all degrees of freedom in $Q^{(1)}$, and no additional terms are compatible in this shot. The weight of $X_1 Y_2 Z_3$ is then updated—and decreases—because it has already been measured. For the next shot, $Q^{(2)}$ begins as the identity and is first aligned with $Z_2 Z_3$, producing $Q^{(2)} = \mathbb{I}_1 Z_2 Z_3$. The next highest-weight compatible term is $Y_1 Z_3$, which updates the remaining identity component of $Q^{(2)}$—which is the one acting on the first qubit—yielding $Q^{(2)} = Y_1 Z_2 Z_3$.

The complete Shadow–Grouping procedure, which integrates the classical shadows framework with grouping techniques, is summarized in Algorithm 2.

Algorithm 2 Shadow–Grouping.

Require: Hamiltonian in Pauli word representation as in (3)
Require: Weight function w
Require: Measurement budget N
for each shot $k = 1, \ldots, N$ **do**
  calculate measurement setting $Q^{(k)}$ using Algorithm 1
  $b \leftarrow$ measure each qubit i in the basis $Q_i^{(k)}$
  for all $O^{(i)}$ QWC with $Q^{(k)}$, estimate $\langle O^{(i)} \rangle$ using (2)
end for
return estimated energy $\langle H \rangle$

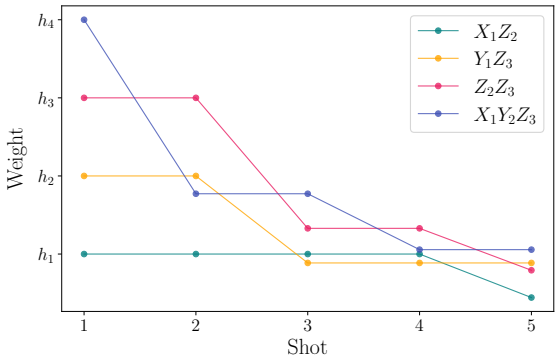

Figure 2: Evolution of weights in Algorithm 1 for the toy Hamiltonian (4) with a measurement budget of $N = 5$ shots.

IV. METHODOLOGY

We evaluated the performance of Algorithms 1 and 2—implemented using the `PennyLane` library [6]—in the context of quantum chemistry. Specifically, we used the Shadow–Grouping protocol to estimate the ground state energies of the hydrogen (H_2) and lithium hydride (LiH) molecules. These molecules serve as standard benchmarks, since quantum chemistry applications require high precision for energy estimates: a prediction is considered chemically accurate when it falls within $\varepsilon_{\text{chem}} = 1.6$ mHa of the true value [4].

The quantum circuit that prepares the ground state of the qubit JW-representations of the corresponding molecular Hamiltonians was obtained from open-access databases [7]. We denote by E_{H_2} and E_{LiH} the exact ground state energies for H_2 and LiH, respectively.

For each molecule, we conducted $L = 100$ independent estimations $\hat{E}^{(1)}, \ldots, \hat{E}^{(L)}$ using the Shadow–Grouping algorithm. Each estimation used a fixed measurement budget of N shots. We then computed the mean estimated energy $\hat{E}$ and the root-mean-square deviation (RMSD) as a measure of statistical error.

V. RESULTS

Hydrogen molecule (H_2): The JW-decomposed Hamiltonian for H_2 is encoded in 4 qubits and contains 15 Pauli terms. As shown in Figure 3, the Shadow–Grouping method consistently outperforms classical shadows across all tested budgets. On average, the energy discrepancy obtained with Shadow–Grouping is 2.6 times smaller than that of classical shadows, with relatively low variability (standard deviation of 0.4). Moreover, Shadow–Grouping achieves chemical accuracy with fewer than 10^5 measurements, whereas classical shadows require two orders of magnitude more.

Lithium hydride molecule (LiH): The Hamiltonian for LiH is encoded in 12 qubits and includes 631 Pauli terms, making it significantly more complex than that of H_2.

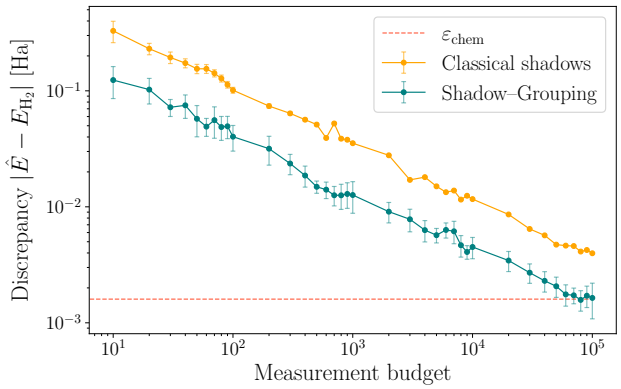

Figure 3: Error decay for H_2 ground state energy estimate.

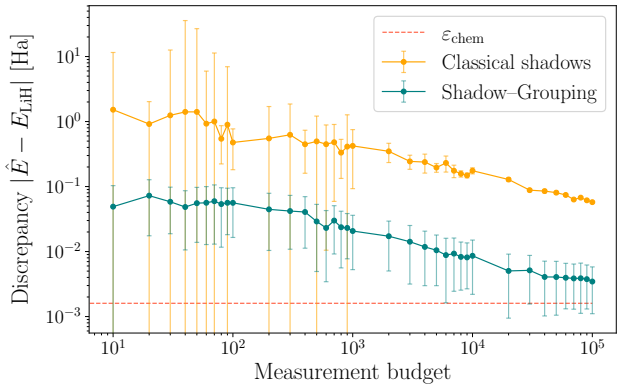

Figure 4: Error decay for LiH ground state energy estimate. Large error bars in classical shadows below 10^3 shots arise from an insufficient sampling of the many terms in the LiH Hamiltonian. Beyond this point, estimations stabilize.

Despite the increased dimensionality, Shadow–Grouping again demonstrates a clear advantage, as shown in Figure 4. Although chemical accuracy is not reached on average, the method consistently achieves it within the uncertainty bounds. Notably, the Shadow–Grouping algorithm is on average 18 ($\pm$5) times more precise than classical shadows and requires approximately three orders of magnitude fewer measurements to reach $\varepsilon_{\text{chem}}$.

These results are summarized in Table I, where we display the number of qubits n and terms M in the JW-encoding of H, the average discrepancy ratio (CS/S–G), and how many shots are needed to achieve chemical accuracy $N(\varepsilon_{\text{chem}})$.

Table I: Performance comparison of classical Shadows (CS) and Shadow–Grouping (S-G) for H_2 and LiH.

Molecule	(n, M)	Method	CS/S–G	$N(\varepsilon_{\text{chem}})$
H_2	(4, 15)	S–G CS	2.6 ± 0.4	$< 10^5$ $> 10^6$
LiH	(12, 631)	S–G CS	18 ± 5	$\sim 10^5$ $\sim 10^8$

Moreover, our method scales effectively with system size: despite the 40-fold increase in Hamiltonian size from H_2 to LiH, Shadow–Grouping outperforms classical shadows. This reveals its potential for larger quantum systems, especially under limited measurement budgets.

VI. CONCLUSIONS

In this thesis, we have explored the intersection of two powerful techniques in quantum information—classical shadows and grouping—and implemented a Shadow–Grouping method, which unifies both approaches and improves measurement efficiency on NISQ devices.

On top of this, through numerical experiments on the H_2 and LiH molecular Hamiltonians, we showed that Shadow–Grouping significantly outperforms standard classical shadows in both accuracy and resource usage. Notably, Shadow–Grouping achieved up to an 18-fold reduction in estimation error and reached chemical precision with two to three orders of magnitude fewer measurements, highlighting its advantage over larger and more complex quantum systems.

Acknowledgments

I want to express my deepest gratitude to Berta Casas, Sergi Masot, and Bruno Juliá-Díaz for their invaluable guidance. I also thank my friends, family, and parents.

[1] H.-Y. Huang, R. Kueng, and J. Preskill. Predicting many properties of a quantum system from very few measurements. *Nature Physics*, 16(10):1050–1057, 2020.

[2] H.-Y. Huang, R. Kueng, and J. Preskill. Efficient estimation of pauli observables by derandomization. *Phys. Rev. Lett.*, 127:030503, Jul 2021.

[3] V. Verteletskyi, T.-C. Yen, and A. F. Izmaylov. Measurement optimization in the variational quantum eigensolver using a minimum clique cover. *The Journal of Chemical Physics*, 152(12):124114, Mar. 2020.

[4] A. Gresch and M. Kliesch. Guaranteed efficient energy estimation of quantum many-body hamiltonians using shadowgrouping. *Nature Communications*, 16:689, 2025.

[5] M. Li, M. Lin, and M. J. S. Beach. Resource-Optimized Grouping Shadow for Efficient Energy Estimation. *Quantum*, 9:1694, 2025.

[6] V. B. et al. Pennylane: Automatic differentiation of hybrid quantum-classical computations. *arXiv preprint arXiv:1811.04968*, 2018.

[7] U. Azad and S. Fomichev. Pennylane quantum chemistry datasets. https://pennylane.ai/datasets/collection/qchem, 2023.

[8] M. H. Bayó. Shadow–Grouping, Git repository. https://github.com/MHBayo/Physics_Thesis.git, 2025.

Millorant les Ombres Clàssiques amb Estratègies d'Agrupació

Author: Marçal Herraiz–Bayó.
Facultat de Física, Universitat de Barcelona, Diagonal 645, 08028 Barcelona, Spain.[*]

Advisors: Berta Casas–Font[1], Sergi Masot–Llima[1]. Tutor: Bruno Juliá–Díaz[2].[†]
(Dated: June 13, 2025)

Resum: Extreure informació dels sistemes quàntics de manera eficient és un dels grans reptes de la computació quàntica. Aquesta tesi se centra en la combinació de dues tècniques complementàries —les ombres clàssiques i l'agrupació d'observables— per millorar l'aprenentatge quàntic. Les ombres clàssiques permeten predir propietats d'un estat quàntic a partir d'un nombre reduït de mesures aleatòries, mentre que les estratègies d'agrupació minimitzen el cost de mesura aprofitant la commutativitat entre observables. S'implementa un enfocament híbrid, anomenat *Shadow—Grouping*, que integra ambdues tècniques per millorar l'eficiència de mesura en els dispositius quàntics actuals. També es mostra la seva eficàcia mitjançant l'estimació de les energies de l'estat fonamental de les molècules d'H_2 i LiH. En comparació amb les ombres clàssiques estàndard, els resultats obtinguts mostren que Shadow–Grouping pot arribar a millorar la precisió fins a un factor de 18, tot requerint diversos ordres de magnitud menys mesures per assolir una precisió química.
Paraules clau: Computació quàntica, ombres clàssiques, agrupació, Hamiltonians de molts cossos.
ODSs: 7, 9, 13.

Objectius de Desenvolupament Sostenible (ODSs o SDGs)

1. Fi de la es desigualtats		10. Reducció de les desigualtats	
2. Fam zero		11. Ciutats i comunitats sostenibles	
3. Salut i benestar		12. Consum i producció responsables	
4. Educació de qualitat		13. Acció climàtica	X
5. Igualtat de gènere		14. Vida submarina	
6. Aigua neta i sanejament		15. Vida terrestre	
7. Energia neta i sostenible	X	16. Pau, justícia i institucions sòlides	
8. Treball digne i creixement econòmic		17. Aliança pels objectius	
9. Indústria, innovació, infraestructures	X		

El contingut d'aquest TFG es relaciona amb l'ODS 7, fita 7.3, per la promoció d'un ús sostenible i eficient dels recursos energètics. També connecta amb l'ODS 9, fita 9.5, per a millorar les capacitats de les tecnologies quàntiques. Finalment, amb l'ODS 13, fita 13.3, per a reduir l'impacte del consum energètic en el canvi climàtic.

GRAPHICAL ABSTRACT