

Active Learning for Quantum Mechanical Measurements

Ruidi Zhu,¹ Ciara Pike-Burke,² and Florian Mintert^{1,3}

¹*Department of Physics, Blackett Laboratory, Imperial College London,
Prince Consort Road, SW7 2AZ, United Kingdom*

²*Department of Mathematics, Huxley Building, Imperial College London,
180 Queen's Gate, SW7 2AZ, United Kingdom*

³*Helmholtz-Zentrum Dresden-Rossendorf, Bautzner Landstraße 400, 01328 Dresden, Germany*

The experimental evaluation of many quantum mechanical quantities requires the estimation of several directly measurable observables, such as local observables. Due to the necessity to repeat experiments on individual quantum systems in order to estimate expectation values of observables, the question of how many repetitions to allocate to a given directly measurable observable arises. We show that an active learning scheme can help to improve such allocations, and the resultant decrease in experimental repetitions required to evaluate a quantity with the desired accuracy increases with the size of the underlying quantum mechanical system.

I. INTRODUCTION

There is quite a discrepancy between the quantum mechanical observables that can be measured *in principle* and those that can be measured *in practice*. Restrictions to local observables or even restrictions to preferred measurement bases are common even in highly controllable experiments with synthetic quantum systems. This results in the necessity to infer quantities that can not be directly measured in terms of observables that can be measured.

While it is well understood how to express a given quantity in terms of expectation values of practically measurable observables [1–4], any such decomposition implicitly assumes expectation values of all involved observables are known. The probabilistic nature of quantum mechanical measurements, however, implies that expectation values can only be estimated with finite accuracy, and any improvement in accuracy requires more repetitions of the same measurement.

The necessity to estimate expectation values of several observables and to increase experimental repetitions in order to improve any such estimate, opens the question of how to allocate experimental resources to the different observables. Lacking a good basis for a different choice, each observable is typically allocated the same number of experimental repetitions [5–7].

The uncertainty sampling strategy [8, 9] in Active Learning (AL) can be used to find better allocations of experimental repetitions. Active learning methods adaptively choose which data to sample in order to minimize the total number of samples required to obtain the desired accuracy. These techniques are typically applied to speed up learning of machine learning models including in speech recognition [10], image retrieval [11], classification tasks [12] or computational physics [13]. In the field of quantum physics active learning is known to be helpful for

the analysis of quantum states [14–17] and optimal control [18]. Active learning can also be regarded as a form of adaptive experimental design, and it is this view of active learning that we will take in this paper.

A key ingredient in the AL scheme is to estimate the informativeness of measurement on each directly measurable observable, and hence to decide which observables to query in consecutive measurements. Concentration inequalities [19–21] provide sound foundations to estimate the uncertainty of variables based on a limited amount of data, especially when this data is generated by some random variables with unknown distributions. Therefore, it is suitable to use concentration inequalities to construct query strategies based on the outcomes of quantum mechanical measurements.

In this paper, we will use concentration inequalities to decide actively which observable to measure in a repetition of an experiment in order to obtain the best-expected improvement of the estimate of a quantity that can not be directly measured. We will show that a dynamical allocation of measurements can help to decrease the total number of repetitions required to estimate a given quantity with desired accuracy, and that this decrease is most pronounced in the regime of a few repetitions.

II. DECISIONS BETWEEN OBSERVABLES

Most quantum mechanical quantities of interest can not be assessed in terms of a single observable. This might be due to practical limitations, as is the case for fidelity with respect to an entangled state or for an entanglement witness: strictly speaking, each of those is a regular observable, but the practical restriction to measurements of single-qubit observables and correlations thereof, implies that several observables need to be measured before the expect-

tation value of the observable of interest can be estimated [2, 22, 23]. This might also be due to more fundamental reasons, as is the case with the gate fidelity [5, 24, 25]: since its experimental evaluation requires the implementation of a gate starting from a complete set of different initial states, estimating a gate fidelity implied performing several independent measurements even if there was not a practical restriction to local observables.

In the following, we will thus consider a general quantum mechanical quantity Q whose experimental evaluation requires estimating the expectations values $\langle S_i \rangle$ of M independent observables S_i . Since Q is a function of the expectation values $\langle S_i \rangle$, the accuracy of the estimate of Q depends on the accuracy with which the expectation values $\langle S_i \rangle$ are estimated. Crucially, the accuracy of each observable does not only depend on the number of measurement repetitions, but it also depends on the actual underlying quantum state [26]. At any given number of repetitions of a σ_z measurement, for example, the accuracy in the estimate of $\langle \sigma_z \rangle$ is higher for a state that is close to a σ_z eigenstate than for a state that would yield more balanced probabilities of the two possible measurement outcomes.

Given an unknown quantum state, one can thus not find an optimal allocation of measurement repetitions to the M different observables to be measured. Only, as data is being taken, one may estimate the accuracy of the different expectation values, and one can use this information in order to decide how to allocate measurement repetitions for subsequent experiments.

In the following, we will thus consider the situation that measurements of several observables S_i that have been performed with n_i repetitions each. Based on the accumulated data, one can estimate the expectation values $\langle S_i \rangle$ and thus the value of the quantity Q of interest with finite accuracy that is limited by the amount of data, *i.e.* the number n_i of measurement repetitions. We will derive a decision rule with AL that helps to identify the observable which, when measured again, will result in the largest available decrease of the inaccuracy in the estimate of Q . With several numerical examples, we will demonstrate that taking data following this AL scheme can help to substantially decrease the number of experiments required to estimate the value of Q with the desired accuracy, and that the gain grows with increasing system size.

III. ACTIVE LEARNING

A. Concentration Inequalities

The law of large numbers [27] states that the expectation value $\langle S \rangle$ of a physical observable S is typically approximated well by the empirical expectation value

$$\langle S \rangle_e = \frac{1}{n} \sum_{i=1}^n s_i, \quad (1)$$

where s_i is the result obtained in the i -th repetition of the measurement of the observable S . The uncertainty in the estimation of $\langle S \rangle$ decreases with the number n of measurement repetitions. This uncertainty can be expressed in terms of concentration inequalities, which states that the upper bound

$$|\langle S \rangle - \langle S \rangle_e| \leq \epsilon(n, \delta) \quad (2)$$

on the deviation between $\langle S \rangle$ and $\langle S \rangle_e$ holds with probability $1 - \delta$.

The explicit form on the upper bound $\epsilon(n, \delta)$ can depend on the underlying problem. The Empirical Bernstein Bound [21, 28, 29]

$$\epsilon_B(n, \delta) = \sqrt{\frac{2v_e \ln \frac{2}{\delta}}{n}} + \frac{7}{3(n-1)} \ln \frac{2}{\delta}, \quad (3)$$

with the empirical variance

$$v_e = \frac{1}{n-1} \sum_{i=1}^n (s_i - \langle S \rangle_e)^2 \quad (4)$$

applies to a wide range of problems. In turn, however, it is not necessarily the best available bound for specific problems. In the case of independent repetitions of a measurement with only two distinct outcomes (a dichotomic observable), the bound

$$\epsilon_D(n, \delta) = \sqrt{\frac{2v \ln \frac{1}{\delta}}{n}} \quad (5)$$

with the actual variance v applies [19]. While this bound generally provides a better estimate of the accuracy of $\langle S \rangle_e$ than the Empirical Bernstein Bound, it has the disadvantage that it is not formulated in terms of the empirical variance v_e , but rather in terms of the actual variance

$$v = p(1-p)(t_1 - t_2)^2, \quad (6)$$

which depends on both the actual probability p to obtain a given outcome and the values t_1 and t_2 that the observable S can adopt (*i.e.* $s_i \in \{t_1, t_2\}$). Since estimating the expectation value of S , or, equivalently, the value of the probability p is the goal of

the experiment, the actual variance v is indeed unknown, so that $\epsilon_D(n, \delta)$ in Eq. (5) is not usable in practice.

A natural remedy seems to replace the actual variance v by its empirical counterpart v_e as defined in Eq. (4). Since, however, in cases with close-to-certain outcomes (*i.e.* $p(1-p) \simeq 0$), the empirical variance v_e tends to be smaller than the actual variance v , this replacement would result in an underestimate of the uncertainty of empirical expectation values of observables with low variance. Any algorithm that is meant to decide to perform measurements of observables with uncertainty estimates would thus decide to perform too many measurements of observables with high variance and too few observables with low variance. In order to find a decision rule that will result in close-to-optimal choices for observables to measure, we aim to find a rule that combines the benefits of being defined in terms of the empirical variance (as in Eq. (3)) with the suitability to dichotomic observables (as in Eq. (5)).

The heuristic ansatz

$$\epsilon_M(n, \delta) = \sqrt{\frac{2v_e}{n} \ln \frac{1}{\delta}} + \frac{(t_1 - t_2)^2 - 4v_e}{4n} \ln \frac{2}{\delta} + \frac{1}{n} \quad (7)$$

includes two additional terms as compared to Eq. (5). With their $1/n$ -dependence, they become negligible in the limit $n \rightarrow \infty$. The last term ensures that $\epsilon_M(n, \delta)$ does not vanish in the case of a few measurements ($n \gtrsim 1$). The second term in Eq. (7) vanishes exactly if v_e adopts its maximal value, and it is the largest for vanishing empirical variance. As such, it results in the desired modification to compensate for the misestimate of low variances. [All explicit examples in the following are based on the numerical value \$\delta = 1/10\$.](#)

B. Uncertainty Reduction

With the ability to estimate the uncertainty of empirical expectation values of directly measurable observables, one can also estimate the uncertainty of the empirical estimate Q_e of the composite quantity of interest. Even though not strictly necessary, we will restrict the following discussion to functions that depend on the observables S_i in a linear fashion, since this is given for quantities like fidelity with respect to a pure state or a unitary gate. Non-linear quantities, such as the von Neumann entropy, would require a generalisation that is feasible (see [Appendix A](#)), but that would make the following discussion unnecessarily technical.

For any given linear function $Q = \sum_i a_i \langle S_i \rangle$, with scalar factors a_i , the empirical estimate of Q reads $Q_e = \sum_i a_i \langle S_i \rangle_e$, and the uncertainty of Q_e can be

estimated with the inequality

$$|Q - Q_e| \leq \sum_{i=1}^N |a_i| \epsilon_i, \quad (8)$$

where ϵ_i is the bound on the uncertainty of $\langle S_i \rangle_e$ following Eq. (7).

The goal at hand is to decrease the inaccuracy of Q_e through identification of the observable to measure that results in the largest possible decrease of the right-hand-side in Eq.(8). To this end, it is desirable to estimate how each of the bounds ϵ_i would change if an additional repetition of the measurement of S_i was performed.

Given the dependence of the bounds in Eq. (7) on the empirical variance, this prediction can be made only approximately. Leaving aside situations with extremely sparse data (*i.e.* $n \gtrsim 1$), the change in the empirical variance following an additional measurement is expected to be negligible; in this approximation, one can thus quantify the expected uncertainty reduction

$$\Delta_i = \epsilon_i(n_i) - \epsilon(n_i + 1), \quad (9)$$

of the estimate of $\langle S_i \rangle_e$, where both $\epsilon_i(n_i)$ and $\epsilon(n_i + 1)$ follow Eq.(7) with the empirical variance σ_e based on n_i measurements.

If all the observable S_i are pairwise non-commuting, then the best available reduction in the uncertainty of Q_e is achieved by measuring the observable S_i that yields the largest value of $|a_i| \Delta_i$. If there are some commuting observables within the set $\{S_i\}$, then it is essential to take into account that commuting observables can be measured in the same run of an experiment. Instead of focusing on individual observables S_i , an algorithm should rather focus on groups G_i of observables, such that all observables in any group do pairwise commute. The expected uncertainty reduction of Q_e upon measurement of the observables in G_i is given by

$$w_i = \sum_{\{j|S_j \in G_i\}} |a_j| \Delta_j, \quad (10)$$

and the group of observables with the largest uncertainty reduction should be measured.

C. Active Learning Algorithm

With the ability to identify the observables to measure that result in the largest uncertainty reduction of the empirical estimate for the quantity Q of interest, we can finally formulate the desired active learning algorithm, which is comprised of the following steps:

- (i) Since no meaningful decision can be taken without any data, **the algorithm is initialised of a single shot of each measurement**. While a single shot is the minimum required to construct an empirical expectation value, at least two shots are required to construct an empirical variance (Eq. (4)). The initialisation will thus include two shots of each of the observables S_i or each of the groups G_i .
- (ii) Once there is enough data to estimate the expected uncertainty reduction w_i (Eq. (10)), the observable S_i or group G_i with the largest expected reduction is selected to be measured in the next step. The outcome of this subsequent measurement is then added to the accumulated data, **and the values of corresponding $\langle S_i \rangle_e$ are updated**. This step is repeated as long as necessary or desired.
- (iii) The process of repeating step (ii) is ended if the empirical estimate Q_e of Q has reached the desired accuracy.

In the examples of explicit implementations of this algorithm discussed below in Sec. IV, this process of estimating Q_e will be compared with a more conventional approach, in which step (ii) is replaced by a selection of observables S_i or group G_i from a fixed list, such that each observable or group is measured approximately as often.

IV. ESTIMATION OF PHYSICAL PROPERTIES WITH ACTIVE LEARNING

This section exemplifies the detailed process of estimating state fidelities and gate fidelities, and the dependence of the benefits of AL on the number of qubits in the underlying systems. All of the subsequent examples are based on numerically simulated measurement outcomes, with the outcomes generated randomly following the quantum mechanical probabilities.

A. State Fidelity

A typical example of a quantity of frequent interest is the fidelity

$$F(\varrho, |\Psi\rangle) = \langle \Psi | \varrho | \Psi \rangle \quad (11)$$

of any given state ϱ with respect to a pure state $|\Psi\rangle$ [2]. In composite quantum systems, it can hardly ever be measured directly, but it can be cast into a

weighted sum of expectation values of directly measurable observables. For any set of mutually orthogonal observables S_i , the state fidelity is of the desired form $F(\varrho, |\Psi\rangle) = \sum_i a_i \langle \Psi | S_i | \Psi \rangle$ with

$$a_i = \frac{\text{tr } \varrho S_i}{\text{tr } S_i^2}. \quad (12)$$

Since in most systems in the context of quantum information processing, the practically accessible observables are restricted to tensor products of Pauli matrices σ_x , σ_y , σ_z and the identity $\mathbf{1}$, the subsequent discussion will assume this choice of observables. Since the identity commutes with all the three Pauli matrices, and measuring an N -qubit observable (*i.e.* a tensor product of N Pauli matrices, but no identity), implies also measuring all observables obtained by replacing Pauli matrices with identities without any additional effort, this situation fits naturally into the setting of commuting observables discussed above.

In order to achieve a sound statistical comparison between fidelity estimates aided by AL and conventional methods, the following discussion is based on state fidelity with respect to states $|\Psi\rangle$ that are randomly chosen from a distribution that is unbiased according to the Haar measure [30] ([see Appendix. B for simulation details](#)). The quantum state ϱ in the state fidelity Eq. (11) is chosen such that the fidelity adopts its maximal value of 1, *i.e.* $\varrho = |\Psi\rangle\langle\Psi|$, but none of the observations made in the following are specific to the case of maximal fidelity.

In particular, in the regime of few measurement repetitions, the data is strongly affected by the statistical fluctuations of measurement results. The empirical estimates of the fidelity will thus typically vary between different realisations of the fidelity estimates. In order to avoid substantial fluctuations in the numerical data, the subsequent discussion will therefore be based on an average over m independent realisations of the same fidelity estimate for any given total number of measurement repetitions (shots) n_T .

With the empirical estimate $F_i(n_T)$ of the fidelity in the i -th realization with a given number of shots n_T , and the exact, theoretically constructed fidelity F , the standard deviation $\sigma(n_T)$ of the fidelity estimate with a given n_T is defined as

$$\sigma(n_T) = \sqrt{\frac{1}{m-1} \sum_{i=1}^m (F - F_i(n_T))^2}. \quad (13)$$

For sufficiently many repeated realisations m , this standard deviation is indeed independent of the statistical fluctuations that are inherent to each individual realisation, and the subsequent examples are based on $m = 10000$ realisations.




FIG. 1: Convergence of fidelity estimates with active learning (triangles) and with the conventionally uniform allocation of repetitions to different observables (squares). Both the standard deviation and the number of shots are depicted on a logarithmic scale. The lines indicate $1/\sqrt{n_T}$ convergence.

Fig. 1 depicts the convergence of such a fidelity estimate for a randomly chosen 4-qubit state with a logarithmic scale for both the number of shots n_T (on the x -axis) and the estimated standard deviation σ (on the y -axis). Data following the estimate aided by AL is depicted with triangles, and data following the conventional estimation strategy is depicted with squares. The black lines depict the $1/\sqrt{n_T}$ dependence that is typical for the reduction of statistical noise. Since $2 \times 3^4 = 162$ shots (*i.e.* two measurement repetitions on each N -qubit observable) are necessary to complete the initialisation stage of the AL algorithm as described in Sec. III C, convergence is shown only for $n_T > 162$. Initially, the convergence with the estimate aided by AL shows a faster decrease than the typical $1/\sqrt{n_T}$ dependence, and it follows the $1/\sqrt{n_T}$ -dependence only after $n_T \simeq 210$ shots. On the other hand, the estimate following the conventional allocation (*i.e.* the total number of shots are evenly distributed to each observable) follows the $1/\sqrt{n_T}$ dependence during the entire process of convergence.

The initial, faster convergence shows that the AL algorithm is indeed capable of identifying the observables to measure that best help to decrease the inaccuracy in the fidelity estimate. Once enough data is accumulated, however, one can decide on an optimal allocation of repetitions to the different observables without accumulating more data. In this case, the adaptive AL method can no longer outperform a strategy with a fixed, but optimised allocation, and the convergence necessarily needs to follow the $1/\sqrt{n_T}$ dependence. Due to the initial, fast convergence, however, the approach aided by AL is expected to outperform conventional approaches also if convergence towards low variances is required, so

that a larger part of the convergence is dominated by the $1/\sqrt{n_T}$ dependence.

The observation that both approaches follow the $1/\sqrt{n_T}$ dependence after sufficiently many shots (for $n_T \gtrsim 210$ in this case), is helpful to define a figure of merit for the improvement of the approach with AL over the conventional approach. The ratio, $n_T^{(c)}/n_T^{(AL)}$, of the number of shots $n_T^{(c)}$ required to achieve a given accuracy of the fidelity with the conventional approach and the number of the shots $n_T^{(AL)}$ to achieve the same accuracy with the approach aided by AL (*i.e.* $\sigma(n_T^{(c)}) = \sigma(n_T^{(AL)})$), is independent of a desired standard deviation of the fidelity estimate as long as this standard deviation is sufficiently small so that the comparison is taken after the initial interval of fast convergence. In the following, we will thus refer to the ratio $n_T^{(c)}/n_T^{(AL)}$ as the *improvement*.

Since state fidelity can be defined for systems with various numbers of qubits, it is well suited to highlight the benefits of AL with increasing system size. The following discussion is thus focused on the state fidelity of an N -qubit system with N ranging from one to six.




FIG. 2: Cumulative distribution of improvements of state fidelity estimations with active learning obtained from statistics with 400 random states for each system size ranging from one to six qubits. While the improvement does depend on the underlying state, there is a clear trend of increasing improvements with a growing qubit number, due to the growing number of observables to choose from.

Fig. 2 depicts the cumulative distribution of the improvement $n_T^{(c)}/n_T^{(AL)}$ found for different system sizes based on fidelity estimates for 400 different random states. In the case of a single qubit (dashed line with circles), one can notice that the improvement is smaller than one in about 25% of the cases. In those cases, the conventional method yields better estimates than the method aided by AL. The dis-

tribution of the observed improvements, however, is skewed towards higher values, and the average improvement does indeed indicate in favour of the AL method.

The only moderate benefit of AL in the estimate of single qubit fidelities can be attributed to the fact that there are only three different directly-measurable observables to choose from. Since, however, the range of different observable settings grows exponentially in the number of qubits, one would expect that the benefits of AL become increasingly pronounced with increasing system size. This expectation is also clearly corroborated by Fig. 2. For three qubits and more, the improvement does always exceed the threshold value of one, and the observed improvements grow steadily with the number of qubits. For $N = 6$ qubits (solid line with downwards triangles), the improvement exceeds the value 1.8 in half of the cases, and the improvement reaches values up to 3; that is, the number of measurements to be taken can be reduced by a factor of 3 without a decrease in the accuracy of the fidelity estimate.

B. Gate Fidelity

The case of state fidelity highlights that the benefits of AL are particularly pronounced if there is a large number of measurement settings to choose from. Since, in the case of gate fidelity, there is a choice for both the initial state and the measurement to be taken on the final state, the estimate aided by AL is potentially particularly beneficial for the estimate of gate fidelities. This section will thus focus on the estimate of gate fidelities. Rather than analysing statistics over randomly chosen gates, this section focuses on the two-qubit controlled-NOT (CNOT) gate and the three-qubit Toffoli gate.

The fidelity of a quantum channel Λ with respect to a gate U for N qubits [24] is given by

$$F(\Lambda, U) = \frac{1}{2^{2N}} \sum_{i,j} \langle i | U^\dagger \Lambda(|i\rangle\langle j|) U |j\rangle, \quad (14)$$

where the summation is performed over two complete sets of orthonormal state vectors.

In order to recast the definition of gate fidelity into an experimentally realisable measurement prescription, it is necessary to expand each of the operators $|i\rangle\langle j|$ in the argument of Λ into a set of actual quantum states. While a set of four quantum states is sufficient for a single qubit, the following analysis is

based on the five states

$$\begin{aligned} |\phi_0\rangle &= |0\rangle, \quad |\phi_1\rangle = |1\rangle \quad \text{and} \\ |\phi_k\rangle &= \frac{1}{\sqrt{2}} \left(|0\rangle + e^{i\frac{2\pi}{3}(k-2)} |1\rangle \right) \quad \text{for } k = 2, 3, 4. \end{aligned} \quad (15)$$

With this choice of states, the gate fidelity for a single qubit can be expressed as

$$F(\Lambda, U) = \frac{1}{4} \sum_{i,j,k} c_{ijk} \langle i | U^\dagger \Lambda(|\phi_k\rangle\langle\phi_k|) U |j\rangle \quad (16)$$

with complex scalar coefficients c_{ijk} . Due to the choice of an over-complete set of states, the values of these coefficients are not uniquely determined, but the choices $c_{000} = 1, c_{111} = 1, c_{01k} = \frac{2}{3}e^{i\frac{2\pi(k-2)}{3}}$ and $c_{10k} = \frac{2}{3}e^{-i\frac{2\pi(k-2)}{3}}$ for $k = 2, 3, 4$, and all remaining terms vanishing is a valid choice.

This generalises straight-forwardly to the gate fidelity for N qubits, with 5^N initial states $|\Phi_k\rangle$ given by tensor products of the single qubits states $|\phi_k\rangle$, and coefficients C_{ijk} given by products of the coefficients c_{ijk} .

An explicit prescription in terms of state preparation, dynamics described by the channel Λ and final measurement is obtained by expanding the operators $\sum_{ij} C_{ijk} U |j\rangle\langle i | U^\dagger$ into the set of observables that can be directly measured. With the set of local Pauli measurements S_i also used in Sec. IV A for the state fidelity, one obtains

$$F(\Lambda, U) = \frac{1}{2^{2N}} \sum_{k=0}^{5^n-1} \sum_{l=0}^{4^n-1} a_{lk} \text{tr}(S_l \Lambda(|\Phi_k\rangle\langle\Phi_k|)), \quad (17)$$

with

$$a_{lk} = \frac{\text{tr}(S_l \sum_{ij} C_{ijk} U |j\rangle\langle i | U^\dagger)}{\text{tr} S_i^2}. \quad (18)$$

The situation regarding the choice of measurements is thus analogous to the case of state fidelity, but in addition to the choice of measurement, there is also the choice of initial state. In every step of the process, the AL algorithm will thus select the most informative initial state and corresponding measurement.

Similarly to the estimate of state fidelities discussed in Sec. IV A, the accuracy of an empirical estimate of the gate fidelity also depends on the actual realisation of random measurement outcomes, and a reliable assessment of the two methods to-be-compared is obtained only in terms of statistics of many independent realisations of the same fidelity estimates.

Fig. 3 depicts the decrease of the standard deviation $\sigma(n_T)$ with the number of shots for the estimation of the perfect fidelity of (a) a two-qubit CNOT




FIG. 3: Convergence comparisons of the fidelity estimations with active learning and conventionally uniform allocations of experimental repetitions for (a) CNOT Gate and (b) Toffoli Gate. Both the standard deviation and the total number of shots are plotted on a logarithmic scale. The reference lines indicate the convergence that scales as $1/\sqrt{n_T}$.

gate, and (b) a three-qubit Toffoli gate, similar to Fig. 1. Triangles denote the case of estimates aided by AL, and squares denote the case in which all measurements are taken with the conventional approach. Qualitatively, the convergence confirms the behavior identified in Fig. 1, but the quantitative details are different: the period of faster convergence in the approach aided by AL lasts until $n_T \approx 2000$ for the CNOT gate and until $n_T \approx 1.1 \times 10^4$ for the Toffoli gate. The improvement as derived from the part of the convergence that satisfies the $1/\sqrt{n_T}$ behaviour is $n_T^{(c)}/n_T^{(AL)} \approx 2$ for the CNOT gate and $n_T^{(c)}/n_T^{(AL)} \approx 2.2$ for the Toffoli gate. With the larger improvement for the CNOT gate and the Toffoli gate as compared to the improvement found for two-qubit and three-qubit state fidelities, Fig. 3 thus confirms the expectations that the benefits of AL are growing with the number of measurement settings to choose from.

V. OUTLOOK

In particular, in the era of *noisy intermediate-scale quantum* (NISQ) devices [31], the estimate of state-fidelities and gate fidelities is a commonly encountered problem [32–34]. The rapidly growing number of observables to be measured makes this an extremely challenging task even for moderate qubit numbers [35, 36]. Due to substantial imperfections, such as drift or uncharacterised cross-talk between qubits in such devices [37–39], there is large uncertainty about a created state or an implemented gate, so a prior allocation of measurement repetitions for the specific state or gate is indeed problematic. The

interactive active learning process for observables to be measured can thus practically facilitate the estimate of fidelities. Since such estimates are at the core of data-driven optimisation processes [40–42], and due to their iterative nature, these optimisations require several fidelity estimates, the proposed algorithm can contribute to our ability to derive practical use from faulty hardware.

The use of the proposed techniques is also not limited to fidelities, but it can also find applicability in *variational quantum algorithms* (VQA) [43–45] in which expectation values of a Hamiltonian or some other operator need to be estimated. The goal of the VQA is the experimental realisation of the quantum state that minimises this expectation value, and iterative optimisation algorithms estimate this expectation value with the same accuracy for all considered states [46]. Since the proposed algorithm does provide not only empirical expectation values, but also bounds on their accuracy, it can also identify the lowest conceivable expectation value at any point in time during the data acquisition. As soon as this value exceeds the expectation value observed with another state, one can safely stop taking data based on this state and start estimating the expectation value with a different state.

With possible extensions to the estimate of quantities like entropy or correlation functions involving products of expectation values, the use of the proposed active learning algorithm has clear potential to become a commonly used tool in the analysis of quantum systems.

ACKNOWLEDGEMENTS

We are indebted to stimulating discussions with Zezhen Wei. Numerical simulations were carried out

on Imperial College High-Performance Computing facilities [47].

-
- [1] T. C. Yen and A. F. Izmaylov, PRX Quantum **2**, 040320 (2021).
- [2] D. F. James, P. G. Kwiat, W. J. Munro, and A. G. White, Phys. Rev. A **64**, 052312 (2001).
- [3] T. C. Yen, V. Verteletskyi, and A. F. Izmaylov, J. Chem. Theory Comput. **16**, 2400 (2020).
- [4] A. F. Izmaylov, T. C. Yen, and I. G. Ryabinkin, Chem. Sci. **10**, 3746 (2019).
- [5] S. Greenaway, F. Sauvage, K. E. Khosla, and F. Mintert, Phys. Rev. Res. **3**, 33031 (2021).
- [6] R. Schmied, J. Mod. Opt. **63**, 1744 (2016).
- [7] A. J. Scott, J. Phys. A **39**, 13507 (2006).
- [8] B. Settles, Computer Sciences Technical Report (2009).
- [9] J. Baldridge and M. Osborne, Proc. of the 2004 Conference on Empirical Methods in Natural Language Processing (2004).
- [10] G. Tur, D. Hakkani-Tür, and R. E. Schapire, Speech Commun. **45**, 171 (2005).
- [11] S. Tong and E. Chang, Proc. of the ACM International Multimedia Conference and Exhibition 10.1145/500141.500159 (2001).
- [12] R. Greiner, A. J. Grove, and D. Roth, Artificial Intelligence **139**, 137 (2002).
- [13] J. Yao, Y. Wu, J. Koo, B. Yan, and H. Zhai, Phys. Rev. Res. **2**, 013287 (2020).
- [14] Y. Ding, J. D. Martín-Guerrero, M. Sanz, R. Magdalena-Benedicto, X. Chen, and E. Solano, Phys. Rev. Lett. **124**, 140504 (2020).
- [15] H. Lange, M. Kebrić, M. Buser, U. Schollwöck, F. Grusdt, and A. Bohrdt, arXiv 10.48550/arxiv.2203.15719 (2022).
- [16] A. Shlosberg, A. J. Jena, P. Mukhopadhyay, J. F. Haase, F. Leditzky, and L. Dellantonio, Quantum **7**, 906 (2023).
- [17] A. A. Melnikov, H. P. Nautrup, M. Krenn, V. Dunjko, M. Tiersch, A. Zeilinger, and H. J. Briegel, Proceedings of the National Academy of Sciences **115**, 1221 (2018), <https://www.pnas.org/doi/pdf/10.1073/pnas.1714936115>.
- [18] Y. Wu, Z. Meng, K. Wen, C. Mi, J. Zhang, and H. Zhai, Chinese Phys. Lett. **37**, 103201 (2020).
- [19] S. Boucheron, G. Lugosi, and O. Bousquet, *Concentration Inequalities* (Springer, Berlin, Heidelberg, 2004).
- [20] S. Zhao, E. Zhou, A. Sabharwal, and S. Ermon, Adv Neural Inf Process Syst **29** (2016).
- [21] V. Mnih, C. Szepesvári, and J. Y. Audibert, Proc. Int. Conf. Mach. Learn., 672 (2008).
- [22] X. Zhang, M. Luo, Z. Wen, Q. Feng, S. Pang, W. Luo, and X. Zhou, Phys. Rev. Lett. **127**, 130503 (2021).
- [23] M. A. and I. N. Chuang, *Quantum Computation and Quantum Information* (Cambridge University Press, 2000).
- [24] A. Gilchrist, N. K. Langford, and M. A. Nielsen, Phys. Rev. A **71**, 062310 (2005).
- [25] M. A. Nielsen, Phys. Lett. A **303**, 249 (2002).
- [26] A. W. Edwards, Nature **186**, 1074 (1960).
- [27] M. J. Evans and J. S. Rosenthal, *Probability and Statistics, the Science of Uncertainty* (W. H. Freeman and Company, 2004).
- [28] A. Maurer and M. Pontil, Proc. of the 22nd Annual Conference on Learning Theory (COLT) (2009).
- [29] P. Shivaswamy and T. Jebara, Proc. of the 13th AIS-TATS **9**, 733 (2010).
- [30] J. Maziero, Braz. J. Phys. **45**, 575 (2015).
- [31] J. Preskill, Quantum **2**, 79 (2018).
- [32] D. Zhu, Z. P. Cian, C. Noel, A. Risinger, D. Biswas, L. Egan, Y. Zhu, A. M. Green, C. H. Alderete, N. H. Nguyen, Q. Wang, A. Maksymov, Y. Nam, M. Cetina, N. M. Linke, M. Hafezi, and C. Monroe, Nat. Commun. **13**, 1 (2022).
- [33] S. Zhang, Y. Lu, K. Zhang, W. Chen, Y. Li, J. N. Zhang, and K. Kim, Nat. Commun. **11**, 1 (2020).
- [34] A. Noiri, K. Takeda, T. Nakajima, T. Kobayashi, A. Sammak, G. Scappucci, and S. Tarucha, Nature **601**, 338 (2022).
- [35] D. Lu, H. Li, D. A. Trottier, J. Li, A. Brodutch, A. P. Krismanich, A. Ghavami, G. I. Dmitrienko, G. Long, J. Baugh, and R. Laflamme, Phys. Rev. Lett. **114**, 10.1103/PHYSREVLETT.114.140505 (2015).
- [36] Y. Zhou, E. M. Stoudenmire, and X. Waintal, Phys. Rev. X. **10**, 10.1103/PhysRevX.10.041038 (2020).
- [37] M. Sarovar, T. Proctor, K. Rudinger, K. Young, E. Nielsen, and R. Blume-Kohout, Quantum **4**, 321 (2020).
- [38] T. Proctor, M. Reville, E. Nielsen, K. Rudinger, D. Lobser, P. Maunz, R. Blume-Kohout, and K. Young, Nature Communications **11**, 5396 (2020).
- [39] C. G. Almudever, L. Lao, X. Fu, N. Khammassi, I. Ashraf, D. Iorga, S. Varsamopoulos, C. Eichler, A. Wallraff, L. Geck, A. Kruth, J. Knoch, H. Bluhm, and K. Bertels, Proc. of the Conference on Design, Automation and Test in Europe (2017).
- [40] B. Dive, A. Pitchford, F. Mintert, and D. Burgarth, Quantum **2**, 80 (2018).
- [41] F. Sauvage and F. Mintert, PRX Quantum **1**, 020322 (2020).
- [42] A. Omran, H. Levine, A. Keesling, G. Semeghini, T. T. Wang, S. Ebadi, H. Bernien, A. S. Zibrov, H. Pichler, S. Choi, J. Cui, M. Rossignolo, P. Rembold, S. Montangero, T. Calarco, M. Endres, M. Greiner, V. Vuletić, and M. D. Lukin, Science **365**, 570 (2019).

- [43] M. Cerezo, A. Arrasmith, R. Babbush, S. C. Benjamin, S. Endo, K. Fujii, J. R. McClean, K. Mitarai, X. Yuan, L. Cincio, and P. J. Coles, *Nat. Rev. Phys.* **3**, 625 (2021).
- [44] D. Wecker, M. B. Hastings, and M. Troyer, *Phys. Rev. A* **92**, 042303 (2015).
- [45] A. Kandala, A. Mezzacapo, K. Temme, M. Takita, M. Brink, J. M. Chow, and J. M. Gambetta, *Nature* **549**, 242 (2017).
- [46] C. N. Self, K. E. Khosla, A. W. R. Smith, F. Sauvage, P. D. Haynes, J. Knolle, F. Mintert, and M. S. Kim, *Npj Quantum Inf.* **7**, 116 (2021).
- [47] Imperial college research computing service.

Appendix A: Uncertainties on Non-linear Quantities

We have restricted our arguments to estimating the values of linear quantities, such as state fidelity and gate fidelity. Non-linear quantities, such as the von Neumann entropy would require a generalisation that is feasible. Here, we present a sketch of this possible generalisation to a quadratic quantity $Q = \sum_i a_i \langle S_i^2 \rangle$ as a concrete example, with scalar factors a_i and expectation values of measurable observables $\langle S_i^2 \rangle$.

The empirical estimate of Q reads $Q_e = \sum_i a_i \langle S_i^2 \rangle_e$. The uncertainty on $\langle S_i^2 \rangle_e$ is bounded by

$$|\langle S_i^2 \rangle - \langle S_i^2 \rangle_e| \leq \epsilon_i |\langle S_i \rangle + \langle S_i \rangle_e|, \quad (\text{A1})$$

where $\epsilon_i = |\langle S_i \rangle - \langle S_i \rangle_e|$ is the bound on the uncertainty of $\langle S_i \rangle_e$ following Eq. (7). The last term in Eq. (A1) is bounded by

$$\begin{aligned} |\langle S_i \rangle_e + \langle S_i \rangle| &= |2\langle S_i \rangle_e + \langle S_i \rangle - \langle S_i \rangle_e| \\ &\leq 2|\langle S_i \rangle_e| + \epsilon_i. \end{aligned} \quad (\text{A2})$$

Therefore, the uncertainty of Q_e can be estimated as

$$|Q - Q_e| \leq \sum_{i=1}^N |a_i| \epsilon_i \left(2|\langle S_i \rangle_e| + \epsilon_i \right). \quad (\text{A3})$$

One can thus use it to calculate the expected uncertainty reduction of Q_e upon measuring each individual observable by replacing ϵ_i appeared in Eq.(7) with $\epsilon_i \left(2|\langle S_i \rangle_e| + \epsilon_i \right)$.

Appendix B: Simulation Details

The construction of random states [30] is comprised of the following steps.

Any N-qubit quantum state $|\psi\rangle$ can be expressed as a superposition of basis states

$$|\psi\rangle = \sum_{k=1}^{2^N} A_k e^{i\theta_k} |\mathbf{c}^{(k)}\rangle, \quad (\text{B1})$$

with phases $\theta_k \in [0, 2\pi)$ and real amplitudes $A_k \in [0, 1]$, satisfying normalisation constraint $\sum_{k=1}^{2^N} A_k^2 = 1$. Here, $\{|\mathbf{c}^{(k)}\rangle\} = \{|c_1^{(k)}, \dots, c_N^{(k)}\rangle\}$ is a reference basis k for the Hilbert space of N qubits.

To generate unbiased random samples of pure states in N-dimensional Hilbert space [30], we first use a uniform random number generator to generate q_1 in the interval $[0, 1]$, and q_k in $[0, 1 - \sum_{j=1}^{k-1} q_j]$ for $k \in \{2, 3, \dots, 2^N - 1\}$, and set $q_N = 1 - \sum_{j=1}^{N-1} q_j$ to ensure the normalisation. Then we shuffle $\sqrt{q_k}$ and randomly assign them to be the real amplitudes A_k for $k \in \{1, 2, \dots, 2^N\}$. To generate random phases, we set the first phase $\theta_1 = 0$, and the rest of phases θ_k randomly generated in $[0, 2\pi)$ for $k = \{2, 3, \dots, 2^N\}$.