Online learning of quantum processes

Asad Raza[1], Matthias C. Caro[1], Jens Eisert[1,2], and Sumeet Khatri[1]

[1]Dahlem Center for Complex Quantum Systems, Freie Universität Berlin, Berlin, Germany
[2]Helmholtz-Zentrum Berlin für Materialien und Energie, Berlin, Germany

Inspired by the expensive resource requirements (both in computation time and number of measurements) of full state and process tomography [1–4], relaxed notions of state and process learning, such as Aaronson's 'pretty good tomography' [5] in the spirit of probably approximately correct (PAC) learning, have recently garnered attention. In this setting, rather than reconstructing the entire density matrix (as in the case of state tomography) or transfer matrix (as in the case of process tomography), the goal is to accurately predict expectation values of the state or process with respect to certain classes of observables. To this end, Ref. [5] showed sample-efficient PAC-learnability of arbitrary quantum states with respect to two-outcome POVMs sampled i.i.d. from any unknown distribution.

In this work, we initiate the study of *online learning* of quantum processes. The most basic formulation of our task is to predict expectation values of the form $\mathrm{Tr}[M_B \mathcal{N}_{A\to B}(\rho_A)]$, where $\mathcal{N}_{A\to B}$ is the unknown quantum channel and (ρ_A, M_B) is a *channel test*, a pair of input state ρ_A and output measurement operator M_B, $0 \le M_B \le \mathbb{1}_B$; see Figure 1. In the PAC-learning setting, the state-measurement pairs are assumed to be given in a batch, and in an i.i.d. manner from some distribution. In the online learning setting that we consider, they are instead given to the learner sequentially and in an adaptive manner, thereby lifting the i.i.d assumption, and the learner aims to achieve minimal regret and a minimal number of mistakes in their predictions. (We define the problem more formally in Section 1 below.) Our work follows prior work on online learning of quantum states [6] and shadow tomography of quantum states [7, 8], in which the observables can also be chosen adaptively.

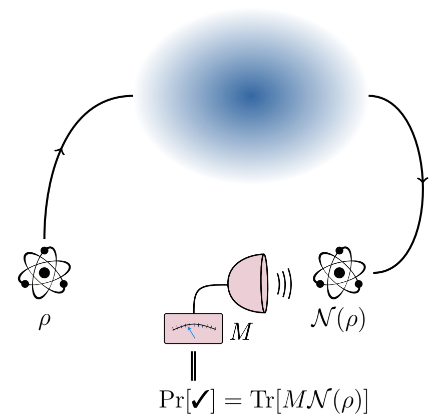

FIGURE 1: To learn about the unknown evolution of a quantum system (symbolized by the blue shaded region and represented mathematically by the quantum channel $\mathcal{N}$), we prepare a probe quantum state ρ, let it evolve, and then measure it according to the POVM $\{M, \mathbb{1} - M\}$.

Contrary to what one might initially anticipate, online learning of quantum processes is not as straightforward as simply online learning the Choi state of the channel. There are two reasons for this. First, the quantities we aim to learn can be written as $\mathrm{Tr}[M_B \mathcal{N}_{A\to B}(\rho_A)] = \mathrm{Tr}[(\rho_A^{\intercal} \otimes M_B) C^{\mathcal{N}}_{A,B}]$, where $C^{\mathcal{N}}_{A,B} = \sum_{i,j=0}^{d_A-1} |i\rangle\langle j|_A \otimes \mathcal{N}_{A\to B}(|i\rangle\langle j|_A)$ is the Choi *matrix* of the channel and d_A is the input dimension of the channel. We can alternatively write $\mathrm{Tr}[M_B \mathcal{N}_{A\to B}(\rho_A)] = d_A \mathrm{Tr}[(\rho_A^{\intercal} \otimes M_B)\Phi^{\mathcal{N}}_{A,B}]$, where now $\Phi^{\mathcal{N}}_{A,B} = \frac{1}{d_A} C^{\mathcal{N}}_{A,B}$ is the Choi state of the channel. The dimension factor resulting from using the Choi state means that good regret and mistake bounds for online learning the Choi state do not give rise to good regret and mistake bounds for online learning the channel. In particular, any attempt to inherit bounds from an online state learning guarantee picks up a multiplicative factor equal to channel's input dimension $d_A = 2^n$, thus leading to exponentially worse bounds for online learning channels. Secondly, online learning algorithms for quantum states (such as those in Ref. [6]) cannot be applied out of the box, due to the partial trace constraint $\mathrm{Tr}_B[\Phi^{\mathcal{N}}_{A,B}] = \mathbb{1}_A/d_A$ on the Choi state. In essence, our goal is to online learn the Choi *matrix* rather than the Choi state. The partial trace condition on the Choi matrix, along with the dimension factor arising from attempting to combine online learning of quantum states with the Choi state, makes the problem that we consider non-trivial.

1 Problem setup

Let $\mathcal{N}_{A\to B}$ be the unknown channel, where A and B are n-qubit systems, and consider an interaction in which a learner faces an adversary for a total of $T \in \mathbb{N}$ time steps. At time step t, the adversary selects a state-measurement pair $(\rho_A^{(t)}, M_B^{(t)})$. The learner generates a hypothesis $\mathcal{N}^{(t)}_{A\to B}$, such that their prediction for the expectation value $\mathrm{Tr}[M_B^{(t)} \mathcal{N}_{A\to B}(\rho_A^{(t)})]$ is $\mathrm{Tr}[M_B^{(t)} \mathcal{N}^{(t)}_{A\to B}(\rho_A^{(t)})]$. The quality of the learner's prediction is measured by a loss function, e.g., $\ell(\mathcal{N}^{(t)}_{A\to B}, \rho_A^{(t)}, M_B^{(t)}) := \left|\mathrm{Tr}[M_B^{(t)} \mathcal{N}_{A\to B}(\rho_A^{(t)})] - \mathrm{Tr}[M_B^{(t)} \mathcal{N}^{(t)}_{A\to B}(\rho_A^{(t)})]\right|$. The setting laid out here, as well as our results, applies to more general tests that involve a reference system R, such that the corresponding expectation value is $\mathrm{Tr}[M_{RB} \mathcal{N}_{A\to B}(\rho_{RA})]$. Our problem setting also generalizes to multi-time processes, in which the evolution can be non-Markovian; see Section 3.

	G-gate channels	Pauli channels
Regret bounds	$\mathcal{O}\left(L\sqrt{TG\log(Gn)}\right)$	$\mathcal{O}\left(L\sqrt{Tn}\right)$
Mistake bounds	$\mathcal{O}\left(\frac{L^2 G\log(Gn)}{\varepsilon^2}\right)$ $\Omega(\min\{2^n, G\})$	$\mathcal{O}\left(\frac{L^2 n}{\varepsilon^2}\right)$ $\Omega(n)$
Computational complexity	$\omega(\text{poly}(n))$ already for $G = \mathcal{O}(n\text{polylog}(n))$	$\tilde{\Omega}(4^n)$

Table 1: **Overview of our main results.** For channels of gate complexity G as well as for n-qubit Pauli channels, we give regret and mistake upper bounds that scale favorably with G and n, respectively. Our bounds hold for general loss functions with Lipschitz constant L. Additionally, complementary mistake lower bounds show that the dependencies on G and n are almost optimal. Finally, we give computational complexity lower bounds for online learning either of the two classes of channels with polynomially many mistakes.

At the end of the T rounds, the learner is evaluated by calculating the total loss and comparing it against a hypothetical scenario in which the learner's predictions could be made after seeing all pairs. This difference is termed 'regret':

$$R_T \coloneqq \sum_{t=1}^{T} \ell(\mathcal{N}^{(t)}, \rho^{(t)}, M^{(t)}) - \inf_{\mathcal{N}} \sum_{t=1}^{T} \ell(\mathcal{N}, \rho^{(t)}, M^{(t)}). \tag{1.1}$$

We aim for online learners that achieve a regret scaling as $\mathcal{O}(\sqrt{T \cdot \text{poly}(n)})$. We also assess the frequency of 'mistakes', defined as events in which the learner's prediction deviates from the correct value by more than a threshold ε, i.e., events in which $\ell(\mathcal{N}^{(t)}, \rho^{(t)}, M^{(t)}) > \varepsilon$. We aim for learners that make no more than $\mathcal{O}(\text{poly}(n))$ ε-mistakes for any (possibly exponential-in-n) number T of time steps.

2 Results

Our results are summarized in Table 1. Ref. [5] left open the question of finding restricted classes of quantum channels for which pretty good process tomography can be realized. We answer this question in a strictly stronger setting, without i.i.d. assumptions, by proving online learnability of two classes of channels. In particular, while it is folklore that general quantum channels are not amenable to online learning (which we explicitly prove), we prove that channels with polynomial gate complexity and Pauli channels are online learnable in the regret- and mistake-bounded models of online learning.

Theorem 1 (Online learning channels of bounded complexity—informal)**.** The class of n-qubit channels that can be implemented by circuits consisting of G arbitrary two-qubit channels can be online learned with regret bound $\mathcal{O}\left(\sqrt{TG\log(Gn)}\right)$ and with ε-mistake bound $\mathcal{O}\left(\frac{G\log(Gn)}{\varepsilon^2}\right)$.

Theorem 2 (Online learning Pauli channels—informal)**.** The class of n-qubit Pauli channels can be online learned with regret bound $\mathcal{O}\left(\sqrt{Tn}\right)$ and with ε-mistake bound $\mathcal{O}\left(\frac{n}{\varepsilon^2}\right)$.

Theorems 1 and 2 give favorably scaling regret and mistake bounds. However, the respective online learning procedures are computationally inefficient. We show that, under reasonable cryptographic assumptions (namely, the hardness of Ring Learning with Errors (RingLWE)), this is unavoidable when aiming for good regret and mistake bounds in these online learning problems:

Theorem 3 (Computational lower bounds for online learning—informal)**.** On the one hand, any online learner that makes at most $\mathcal{O}(\text{poly}(n))$ many (1/3)-mistakes in online learning n-qubit Pauli channels has to use runtime exponential in n. On the other hand, assuming that RingLWE cannot be solved by classical polynomial-time algorithms, then already for $G = \mathcal{O}(n\text{polylog}(n))$ there is no polynomial-time online learner that makes at most $\mathcal{O}(\text{poly}(n))$ many (1/3)-mistakes in online learning G-gate n-qubit quantum channels.

3 Extensions and applications

Our results extend to the online learning of probabilistic mixtures of any exponentially large set of known channels.

Theorem 4 (Online learning convex mixtures of known channels—informal)**.** Given an arbitrary set of $K > 0$ known and fixed quantum channels, and convex mixture of these channels can be online learned with regret bound $\mathcal{O}\left(\sqrt{T \log K}\right)$ and ε-mistake bound $\mathcal{O}\left(\frac{\log(K)}{\varepsilon^2}\right)$.

As an application of our proof technique in Theorem 2, we connect the task of shadow tomography of Pauli channels to classical adaptive data analysis [9]. In the task of shadow tomography, the goal is to provide ε-accurate estimates for all m given channel tests (which could be chosen adversarially) with probability at least $1 - \delta$. By sampling from the Pauli error rate distribution using Bell basis measurements on the Choi state, and invoking known guarantees from classical adaptive data analysis [10], we prove the following result.

Theorem 5 (Shadow tomography of Pauli channels—informal)**.** Shadow tomography of an arbitrary n-qubit Pauli channel can be solved using

$$k = \mathcal{O}\left(\frac{\sqrt{n}\log(m)\log^{3/2}((\varepsilon\delta)^{-1})}{\varepsilon^3}\right) \tag{3.1}$$

copies of the unknown Pauli channel.

Multi-time processes. Finally, our results go beyond quantum channels to multi-time processes [11, 12]. Notably, such processes can be non-Markovian, maintaining a memory over time; see the blue process in Figure 2. Here, we make use of the theory of quantum combs/quantum strategies [13, 14], and the fact that the Born rule generalizes to multi-time process [13–15]. With this, online learning of multi-time processes is formulated analogously to channels as the task of predicting quantities of the form $\mathrm{Tr}[EN]$, where N is the Choi representation of the unknown process and E is a tester, representing the red process in Figure 2. We show that our regret and mistake bounds from Theorem 1 and Theorem 4 also generalize to (non-Markovian) multi-time quantum processes with analogous structure: (a) implementable with polynomially many multi-time processes and (b) convex combinations of (exponentially many) known multi-time processes. We also provide a shadow tomography result for multi-time process that is analogous to Theorem 5.

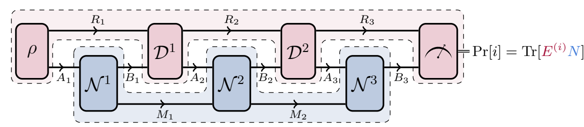

Figure 2: Tests for a multi-time process (depicted in red) consist of an input state ρ, generally entangled with a reference system, along with adaptive operations $\mathcal{D}^k$, followed by a joint system-reference measurement. The measurement probabilities are given by a generalization of the Born rule.

4 Outlook

Our work initiates the study of quantum channel learning in the framework of online learning, which in particular incorporates learning in an environment that presents adaptively chosen challenges. In addition to formalizing this framework, our main contributions are as follows: On the one hand, we identified two natural classes of channels – channels of bounded gate complexity and Pauli channels – that are learnable with favorably scaling regret and mistake bounds. These results demonstrate that two classes of physically relevant noise models can be learned even in a setting with adversarial challenges. On the other hand, we established fundamental obstacles to computationally efficient online learning even for these restricted classes of channels. This highlights that even when online learning is possible in principle, it may remain computationally infeasible.

A natural next step is to find other interesting classes of channels that admit at most polynomial mistake bounds, ideally using only computationally efficient online learners. Given the important role of online learning quantum states in state tomography procedures, we envision positive results, both ours and future ones, on online learning restricted classes of channels to serve as a stepping stone towards shadow tomography procedures for such classes.

Finally, our results may be the starting point of a program of "onlinification" of quantum learning in other scenarios. For instance, inspired by recent work [16, 17], one may attempt to circumvent exponential lower bounds in online learning arbitrarily complex channels by imposing restrictions on the behavior of the adversary, for instance, with respect to the challenges that they can pose. And given recent progress in questions of Hamiltonian learning [18–23], online learning variants of the standard Hamiltonian learning task could give insights into whether one can learn to fine-tune a Hamiltonian evolution according to adaptive feedback.

References

[1] M. Mohseni, A. T. Rezakhani, and D. A. Lidar, "Quantum-process tomography: Resource analysis of different strategies", Physical Review A **77**, 032322 (2008), arXiv:quant-ph/0702131.

[2] J. Haah, A. W. Harrow, Z. Ji, X. Wu, and N. Yu, "Sample-optimal tomography of quantum states", IEEE Transactions on Information Theory **63**, 5628 (2017), arXiv:1508.01797.

[3] R. O'Donnell and J. Wright, "Efficient quantum tomography", in Proceedings of the forty-eighth annual ACM symposium on Theory of Computing (2016), pp. 899–912.

[4] A. Oufkir, *Sample-optimal quantum process tomography with non-adaptive incoherent measurements*, 2023, arXiv:2301.12925.

[5] S. Aaronson, "The learnability of quantum states", Proceedings of the Royal Society A: Mathematical, Physical and Engineering Sciences **463**, 3089 (2007), arXiv:quant-ph/0608142.

[6] S. Aaronson, X. Chen, E. Hazan, S. Kale, and A. Nayak, "Online learning of quantum states", Journal of Statistical Mechanics: Theory and Experiment **2019**, 124019 (2019), arXiv:1802.09025.

[7] S. Aaronson, "Shadow tomography of quantum states", in Proceedings of the 50th Annual ACM SIGACT Symposium on Theory of Computing, STOC 2018 (2018), pp. 325–338, arXiv:1711.01053.

[8] C. Bădescu and R. O'Donnell, "Improved quantum data analysis", TheoretiCS **Volume 3**, 10.46298/theoretics.24.7 (2024), arXiv:2011.10908.

[9] C. Dwork, V. Feldman, M. Hardt, T. Pitassi, O. Reingold, and A. Roth, "Preserving Statistical Validity in Adaptive Data Analysis", in Proceedings of the Forty-Seventh Annual ACM Symposium on Theory of Computing, STOC 2015 (2015), pp. 117–126, arXiv:1411.2664.

[10] R. Bassily, K. Nissim, A. Smith, T. Steinke, U. Stemmer, and J. Ullman, "Algorithmic Stability for Adaptive Data Analysis", SIAM Journal on Computing **50**, STOC16–377-STOC16 (2016), arXiv:1511.02513.

[11] F. A. Pollock, C. Rodríguez-Rosario, T. Frauenheim, M. Paternostro, and K. Modi, "Non-Markovian quantum processes: Complete framework and efficient characterization", Physical Review A **97**, 012127 (2018), arXiv:1512.00589.

[12] S. Milz and K. Modi, "Quantum stochastic processes and quantum non-Markovian phenomena", PRX Quantum **2**, 030201 (2021), arXiv:2012.01894.

[13] G. Gutoski and J. Watrous, "Toward a general theory of quantum games", in Proceedings of the Thirty-Ninth Annual ACM Symposium on Theory of Computing, STOC '07 (2007), pp. 565–574, arXiv:quant-ph/0611234.

[14] G. Chiribella, G. M. D'Ariano, and P. Perinotti, "Theoretical framework for quantum networks", Physical Review A **80**, 022339 (2009).

[15] M. Ziman, "Process positive-operator-valued measure: A mathematical framework for the description of process tomography experiments", Physical Review A **77**, 062112 (2008), arXiv:0802.3862.

[16] H.-Y. Huang, S. Chen, and J. Preskill, "Learning to predict arbitrary quantum processes", arXiv:2210.14894 (2022).

[17] M. C. Caro, "Learning quantum processes and Hamiltonians via the Pauli transfer matrix", arXiv:2212.04471 (2022).

[18] J. Haah, R. Kothari, and E. Tang, "Optimal learning of quantum Hamiltonians from high-temperature Gibbs states", in 2022 IEEE 63rd Annual Symposium on Foundations of Computer Science (FOCS) (IEEE, 2022), pp. 135–146.

[19] D. Stilck França, L. A. Markovich, V. V. Dobrovitski, A. H. Werner, and J. Borregaard, "Efficient and robust estimation of many-qubit Hamiltonians", arXiv:2205.09567 (2022).

[20] F. Wilde, A. Kshetrimayum, I. Roth, D. Hangleiter, R. Sweke, and J. Eisert, "Scalably learning quantum many-body Hamiltonians from dynamical data", arXiv:2209.14328 (2022).

[21] H.-Y. Huang, Y. Tong, D. Fang, and Y. Su, "Learning many-body Hamiltonians with Heisenberg-limited scaling", Physical Review Letters **130**, 200403 (2023).

[22] A. Bakshi, A. Liu, A. Moitra, and E. Tang, "Learning quantum Hamiltonians at any temperature in polynomial time", arXiv:2310.02243 (2023).

[23] A. Bakshi, A. Liu, A. Moitra, and E. Tang, "Structure learning of Hamiltonians from real-time evolution", arXiv preprint arXiv:2405.00082 (2024).