

Improved regret bounds for structured online learning of quantum states

Akshay Bansal*

Jiahui Liu[†]

Abstract

Shadow tomography is fundamental to quantum information processing that focuses on predicting measurement outcomes rather than reconstructing the full state. Its online variant models adaptive and potentially adversarial measurement scenarios, where a learner sequentially predicts outcomes while competing with the best fixed quantum state in hindsight.

We show that exploiting additional structure in the measurements leads to significantly stronger regret guarantees. In particular, under the assumption that the adversarial measurements have bounded Frobenius norm, we analyze Projected Online Gradient Descent and derive regret bounds that depend on intrinsic structural properties, such as rank or sparsity, rather than the ambient Hilbert space dimension. As a complementary result, we show that one can achieve logarithmic regret, independent of both the number of qubits and measurement outcomes, for multi-outcome measurements under squared L_2 loss. These results demonstrate that incorporating realistic structural assumptions can substantially enhance the learnability of quantum states in online environments.

1 Introduction

Shadow tomography is a fundamental task introduced in [Aar18] that is essential for both theoretical research and practical applications in a number of disciplines. It aims at extracting specific properties of the target state. Specifically, given multiple copies of an unknown quantum state and a set of two-outcome measurements (known apriori), the goal is to produce a representative quantum state such that the measurement outcome distributions for the representative state closely approximate those of the unknown quantum state. However, it has practical limitations. Firstly, it does not account for adversarial environments or those that evolve over time as all the measurements are considered to be drawn independently from a fixed distribution. Furthermore, not all quantum measurements are accessible or feasible in a given experimental setup, as their feasibility is dictated by physical constraints. To address these considerations, [ACH⁺18] formalized a notion that incorporates real-time measurement constraints into tomography, framing the problem as an instance of online learning of quantum states where they analyzed standard algorithms such as Regularized Follow-the-Leader (RFTL) and the Matrix Multiplicative Weights (MMW) method.

Motivation for exploiting structure in quantum learning While [ACH⁺18] considered a much general setting of online learning of quantum states, many practical applications admit additional structural properties that can be exploited to obtain improved regret bounds in terms of the number of qubits and time horizon. We focus on two sources of structure that are not only relevant to practical quantum computing but can also be leveraged beneficially in the regret bound analysis.

*Technische Universität Wien, Austria. akshay.bansal@tuwien.ac.at

[†]Fujitsu Research of America, USA. jliu@fujitsu.com

The first concerns the measurement operators themselves. In many applications, the relevant effects are not arbitrary full-dimensional operators; instead, they may be low-rank, sparse in a preferred basis, or supported on a physically meaningful truncated subspace. For example, it is now well established that *optimal* full-state tomography can be achieved using symmetric informationally complete (SIC) POVMs [OP15] (rank-one operators), and this has also been demonstrated experimentally on an ion-trap quantum processor [SMP⁺22]. Similarly, structured truncations of parity- and syndrome-type measurements arise in quantum error correction settings [LRBL⁺23, MC24]. More broadly, Hilbert–Schmidt/Frobenius geometry also plays a role in compressed-sensing tomography (with sparse effects in Pauli basis) and distinguishability problems [GLF⁺10, FGLE12, Spe14].

A second source of exploitable structure in quantum learning arises from the choice of loss function. Previous works consider a more general class of loss functions that are assumed only to be convex and Lipschitz continuous. In contrast, many practical applications employ metric losses, such as the ℓ_1 and ℓ_2 distances, to quantify the discrepancy between the learner’s predicted outcome and the actual measurement outcome [WCS⁺21, KSPM25]. These losses naturally interact with vector-valued predictions in the k -outcome setting and can be exploited to obtain sharper regret guarantees.

Motivated by the high-level objective of exploiting structural properties to improve the learnability of quantum states in the online setting, we investigate whether additional assumptions on the measurement operators or the loss functions can lead to stronger regret guarantees. In particular, we address the following research questions:

1. Can improved regret bounds be obtained when the adversarial measurement operators possess specific structural properties, such as sparsity or low rank?
2. By focusing on particular distance-based loss functions, can sharper regret bounds be achieved using existing algorithms for commonly employed metrics such as the L_1 or L_2 distances?

While the first question is the primary focus of this work, we address the second only briefly and leave a more general treatment to an ongoing research ¹.

2 Our Results

Our first result concerns the setting in which the adversarial measurement operators admit additional structure. We analyze the Projected Online Gradient Descent (OGD) algorithm [Haz16, Chapter 3] for the problem of online learning of quantum states under the assumption that the measurement operators have bounded Frobenius norm (denoted by $\|\cdot\|_F$). Our analysis yields the following informal guarantee.

Theorem 1 (Informal). *Let T denote the time horizon and let n be the number of qubits. Suppose that at each round $t \in [T]$, the loss function is convex and L -Lipschitz. Further assume that the adversarial measurement operators satisfy $\|E_t\|_F \leq B$ for all $t \in [T]$. Then the regret satisfies*

$$R_T \leq \mathcal{O}(LB\sqrt{T}).$$

¹The L_1 and L_2 losses considered here are only representative examples of the loss functions commonly studied in online learning. In future work, it would be interesting to extend our analysis to a broader class of structured loss functions, including strictly convex losses, kernel-based losses, exp-concave losses, and other loss functions frequently encountered in the online learning literature.

We next examine the computational complexity of the successive updates produced by Projected Online Gradient Descent (OGD). In particular, we show that the cost of each OGD update is no worse than that of the Regularized-Follow-the-Leader (RFTL) updates employed in the regret analysis of [ACH⁺18]. This demonstrates that, even in the quantum setting, the iterates of OGD can be computed efficiently, making the algorithm computationally attractive in addition to its favorable regret guarantees. More formally, we have the following observation.

Observation 1. *The classical computational complexity of each successive guesstimate update produced by Projected Online Gradient Descent is $\mathcal{O}(d^\omega)$, where $2 \leq \omega \leq 2.37$ denotes the matrix multiplication exponent and $d \times d$ is the dimension (size) of the effects.*

Combining Theorem 1 with the general regret bound of [ACH⁺18] yields the following immediate consequences.

Corollary 1. *If the adversarial effects satisfy $\text{rank}(E_t) \leq r$ for all $t \in [T]$, then*

$$R_T \leq \mathcal{O}(L\sqrt{\min\{r, n\} T}).$$

In an earlier work by [CW20], a similar regret bound for low-rank effects was discovered by analyzing Regularized-Follow-the-Leader (RFTL) method with Tallis-2 entropy as a regularizer, but this work of ours largely generalizes it to include various other structured effects.

Corollary 2. *If the adversarial effects are κ -sparse for all $t \in [T]$, then*

$$R_T \leq \mathcal{O}(L\sqrt{\min\{\kappa, n\} T}).$$

In particular, these bounds recover the previously known guarantees for general measurements, while yielding potentially sharper rates when the measurement operators are low rank or sparse. Such structured measurements arise naturally in many quantum information settings, as discussed in Section 1.

Our follow-up result focuses on the multi-outcome measurement setting and leverages a specific structure in the loss function. We consider the Follow-the-Leader (RFTL) algorithm under the squared L_2 -norm loss to obtain a logarithmic regret bound.

Proposition 1 (Informal). *Let T be the time horizon. At each round $t \in [T]$, the adversary selects a K -outcome measurement $(E_{t,1}, E_{t,2}, \dots, E_{t,K})$. The learner incurs the squared L_2 loss*

$$\|x_t - b_t\|_2^2,$$

where

$$b_t = \begin{bmatrix} \text{tr}(E_{t,1}\rho) \\ \text{tr}(E_{t,2}\rho) \\ \vdots \\ \text{tr}(E_{t,K}\rho) \end{bmatrix}$$

denotes the vector of true measurement outcome probabilities for an unknown quantum state ρ and $x_t \in [0, 1]^K$ is the learner's prediction. Then the regret of the RFTL algorithm is bounded by

$$R_T \leq \mathcal{O}(\log T).$$

This establishes logarithmic regret in the time horizon for the squared L_2 loss under K -outcome measurements which is independent of the number of the qubits or the number of possible measurement outcomes, demonstrating that additional structure in the loss function can lead to substantially improved guarantees.

References

- [Aar18] Scott Aaronson. Shadow tomography of quantum states. In *Proceedings of the 50th Annual ACM SIGACT Symposium on Theory of Computing (STOC)*, page 325–338, 2018.
- [ACH⁺18] Scott Aaronson, Xinyi Chen, Elad Hazan, Satyen Kale, and Ashwin Nayak. Online learning of quantum states. In *Advances in Neural Information Processing Systems*, volume 31, 2018.
- [CW20] Yifang Chen and Xin Wang. More practical and adaptive algorithms for online quantum state learning (2020). *arXiv preprint arXiv:2006.01013*, 2020.
- [FGLE12] Steven T Flammia, David Gross, Yi-Kai Liu, and Jens Eisert. Quantum Tomography via Compressed Sensing: Error Bounds, Sample Complexity and Efficient Estimators. *New Journal of Physics*, 14(9):095022, 2012.
- [GLF⁺10] David Gross, Yi-Kai Liu, Steven T. Flammia, Stephen Becker, and Jens Eisert. Quantum state tomography via compressed sensing. *Phys. Rev. Lett.*, 105:150401, 2010.
- [Haz16] Elad Hazan. Introduction to online convex optimization. *Foundations and Trends in Optimization*, 2(3-4):157–325, 2016.
- [KSPM25] Viacheslav Kuzmin, Wilfrid Somogyi, Ekaterina Pankovets, and Alexey Melnikov. Method for noise-induced regularization in quantum neural networks. *Advanced Quantum Technologies*, 8(12):e00603, 2025.
- [LRBL⁺23] Francois-Marie Le Régent, Camille Berdou, Zaki Leghtas, Jérémie Guillaud, and Mazhar Mirrahimi. High-performance repetition cat code using fast noisy operations. *Quantum*, 7:1198, 2023.
- [MC24] Z. M. McIntyre and W. A. Coish. Flying-cat parity checks for quantum error correction. *Phys. Rev. Res.*, 6:023247, 2024.
- [OP15] Hiromichi Ohno and Dénes Petz. Examples of conditional SIC-POVMs. *Quantum Information Processing*, 14(10):3983–3999, 2015.
- [SMP⁺22] Roman Stricker, Michael Meth, Lukas Postler, Claire Edmunds, Chris Ferrie, Rainer Blatt, Philipp Schindler, Thomas Monz, Richard Kueng, and Martin Ringbauer. Experimental single-setting quantum state tomography. *PRX Quantum*, 3:040310, 2022.
- [Spe14] Dominique Spehner. Quantum correlations and distinguishability of quantum states. *Journal of Mathematical Physics*, 55(7):075211, 2014.
- [WCS⁺21] Zhen Wang, Yanzhu Chen, Zixuan Song, Dayue Qin, Hekang Li, Qiujiang Guo, H. Wang, Chao Song, and Ying Li. Scalable evaluation of quantum-circuit error loss using clifford sampling. *Phys. Rev. Lett.*, 126:080501, 2021.