Proceedings of Machine Learning Research vol 201:1–3, 2023 34th International Conference on Algorithmic Learning Theory

# Online Self-Concordant and Relatively Smooth Minimization, With Applications to Online Portfolio Selection and Learning Quantum States

**Chung-En Tsai** CHUNGENTSAI1007@GMAIL.COM
*Department of Computer Science and Information Engineering, National Taiwan University*

**Hao-Chung Cheng** HAOCHUNG.CH@GMAIL.COM
*Department of Electrical Engineering and Graduate Institute of Communication Engineering, National Taiwan University*
*Department of Mathematics, National Taiwan University*
*Hon Hai (Foxconn) Quantum Computing Centre*
*Center for Quantum Science and Engineering, National Taiwan University*
*Physics Division, National Center for Theoretical Sciences, Taipei 10617, Taiwan*

**Yen-Huan Li** YENHUAN.LI@CSIE.NTU.EDU.TW
*Department of Computer Science and Information Engineering, National Taiwan University*
*Department of Mathematics, National Taiwan University*
*Center for Quantum Science and Engineering, National Taiwan University*

**Editors:** Shipra Agrawal and Francesco Orabona


## Abstract

This work considers online portfolio selection (OPS) and online learning quantum states with the logarithmic loss. The problem of designing an optimal OPS algorithm, in both regret and efficiency, has been open for more than 30 years (Cover, 1991; Cover and Ordentlich, 1996; Helmbold et al., 1998; Nesterov, 2011; Orseau et al., 2017; Luo et al., 2018; van Erven et al., 2020; Mhammedi and Rakhlin, 2022; Zimmert et al., 2022). Online learning quantum states is a generalization of OPS to the quantum setup (Lin et al., 2021; Zimmert et al., 2022). The dimension of a quantum state grows exponentially with the number of qubits, so scalability with respect to the dimension becomes a critical issue in the quantum setup.

We formulate the two problems as online convex optimization where the loss functions are self-concordant barriers and smooth relative to a convex function $h$. We analyze the regret of online mirror descent with $h$ as the regularizer. Then, based on the analysis, we prove the following in a *unified* manner. Denote by $T$ the time horizon and $d$ the parameter dimension.

1. For OPS, the regret of $\widetilde{\text{EG}}$, a variant of exponentiated gradient due to Helmbold et al. (1998), is $\tilde{O}(T^{2/3}d^{1/3})$. This improves on the original $\tilde{O}(T^{3/4}d^{1/2})$ regret bound for $\widetilde{\text{EG}}$.
2. For OPS, the regret of LB-OMD (online mirror descent with the logarithmic barrier) is $\tilde{O}(\sqrt{Td})$. The regret bound and the per-iteration time are the same as those of Soft-Bayes due to Orseau et al. (2017) up to logarithmic terms.
3. For online learning quantum states with the logarithmic loss, the regret of online mirror descent with the log-determinant function is also $\tilde{O}(\sqrt{Td})$. Its per-iteration time is $O(d^3)$, shorter than all existing algorithms we know.


† Extended abstract. Full version appears as (Tsai et al., 2022).

Figure 1 presents the current efficiency-regret Pareto frontier for OPS, where the two results in bold are our contributions.

In contrast to existing works on online convex optimization (Nesterov, 2011; Abernethy et al., 2012; van Erven et al., 2020), our analysis does not assume $h$ to be a self-concordant barrier. Instead, we exploit the self-concordant barrier property of the *loss functions*, and only require $h$ to satisfy the relative smooth property. One advantage of our approach is that it includes $\widetilde{\text{EG}}$ as an instance, though the entropic regularizer is not a barrier function.

**Figure 1:** Current efficiency-regret Pareto frontier for OPS, assuming $T \geq 4d/\log d$. The figure is modified from the one by Zimmert et al. (2022).

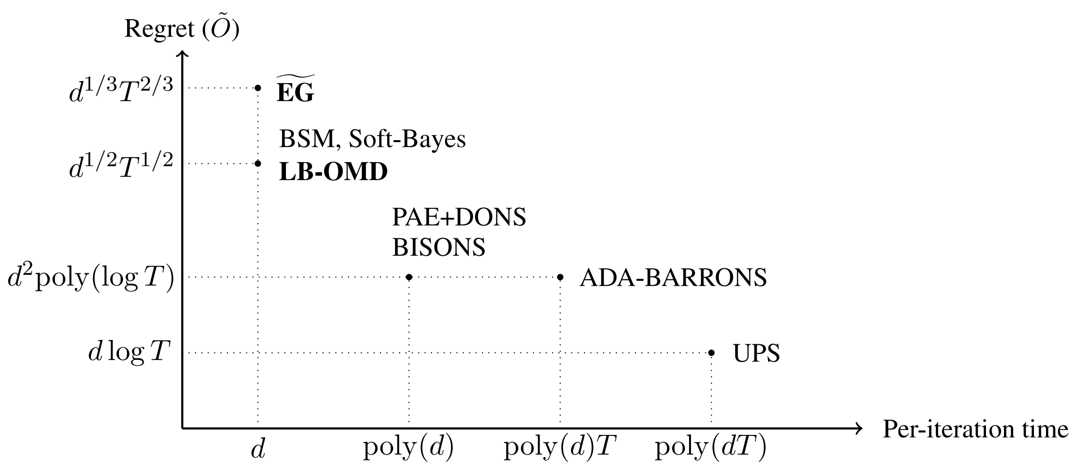


**Keywords:** online convex optimization, online portfolio selection, learning quantum states, online mirror descent, self-concordant barrier, relative smoothness, logarithmic loss

## Acknowledgments

We thank the anonymous reviewers for their inspiring comments.

C.-E. Tsai and Y.-H. Li are supported by the Young Scholar Fellowship (Einstein Program) of the National Science and Technology Council of Taiwan under grant numbers MOST 108-2636-E-002-014, MOST 109-2636-E-002-025, MOST 110-2636-E-002-012, MOST 111-2636-E-002-019, and NSTC 112-2636-E-002-003 and by the research project "Pioneering Research in Forefront Quantum Computing, Learning and Engineering" of National Taiwan University under grant numbers NTU-CC-111L894606 and NTU-CC-112L893406.

H.-C. Cheng is supported by the Young Scholar Fellowship (Einstein Program) of the National Science and Technology Council (NSTC) in Taiwan (R.O.C.) under Grant MOST 111-2636-E-002-026, Grant MOST 111-2119-M-007-006, Grant MOST 111-2119-M-001-004, and is supported by the Yushan Young Scholar Program of the Ministry of Education in Taiwan (R.O.C.) under Grant NTU-111V1904-3, and Grant NTU-111L3401, and by the research project "Pioneering Research in Forefront Quantum Computing, Learning and Engineering" of National Taiwan University under Grant No. NTU-CC-111L894605."

## References


Jacob D. Abernethy, Elad Hazan, and Alexander Rakhlin. Interior-Point Methods for Full-Information and Bandit Online Learning. *IEEE Trans. Inf. Theory*, 58(7):4164–4175, 2012.

Thomas M. Cover. Universal portfolios. *Math. Financ.*, 1(1):1–29, 1991.

Thomas M. Cover and Erick Ordentlich. Universal portfolios with side information. *IEEE Trans. Inf. Theory*, 42(2):348–363, 1996.

David P. Helmbold, Robert E. Shapire, Yoram Singer, and Manfred K. Warmuth. On-line portfolio selection using multiplicative updates. *Math. Financ.*, 8(4):325–347, 1998.

Chien-Ming Lin, Yu-Ming Hsu, and Yen-Huan Li. An Online Algorithm for Maximum-Likelihood Quantum State Tomography. In *24th Annu. Conf. Quantum Information Processing*, 2021. arXiv:2012.15498.

Haipeng Luo, Chen-Yu Wei, and Kai Zheng. Efficient online portfolio with logarithmic regret. In *Adv. Neural Information Processing Systems 31*, 2018.

Zakaria Mhammedi and Alexander Rakhlin. Damped Online Newton Step for Portfolio Selection. In *Proc. 35th Annu. Conf. Learning Theory*, pages 5561–5595, 2022.

Yurii Nesterov. Barrier subgradient method. *Math. Program., Ser. B*, 127:31–56, 2011.

Laurent Orseau, Tor Lattimore, and Shane Legg. Soft-Bayes: Prod for mixtures of experts with log-loss. In *Proc. 28th Int. Conf. Algorithmic Learning Theory*, pages 372–399, 2017.

Chung-En Tsai, Hao-Chung Cheng, and Yen-Huan Li. Online Self-Concordant and Relatively Smooth Minimization, With Applications to Online Portfolio Selection and Learning Quantum States. 2022. arXiv:2210.00997.

Tim van Erven, Dirk van der Hoeven, Wojciech Kotłowski, and Wouter M. Koolen. Open Problem: Fast and Optimal Online Portfolio Selection. In *Proc. 33rd Annu. Conf. Learning Theory*, pages 3864–3869, 2020.

Julian Zimmert, Naman Agarwal, and Satyen Kale. Pushing the efficiency-regret Pareto frontier for online learning of portfolios and quantum states. In *Proc. 35th Annu. Conf. Learning Theory*, pages 182–226, 2022.