
Maximum-Likelihood Quantum State Tomography by Soft-Bayes

Anonymous Author(s)

Affiliation

Address

email

Abstract

1 Quantum state tomography (QST), the task of estimating an unknown quantum state
2 given measurement outcomes, is essential to building reliable quantum computing
3 devices. Whereas computing the maximum-likelihood (ML) estimate corresponds
4 to solving a finite-sum convex optimization problem, the objective function is not
5 smooth nor Lipschitz, so most existing convex optimization methods lack sample
6 complexity guarantees; moreover, both the sample size and dimension grow expo-
7 nentially with the number of qubits in a QST experiment, so a desired algorithm
8 should be highly scalable with respect to the dimension and sample size, just like
9 stochastic gradient descent. In this paper, we propose a stochastic first-order algo-
10 rithm that computes an ε -approximate ML estimate in $O((D \log D)/\varepsilon^2)$ iterations
11 with $O(D^3)$ per-iteration time complexity, where D denotes the dimension of the
12 unknown quantum state and ε denotes the optimization error. Our algorithm is an
13 extension of Soft-Bayes to the quantum setup.

14 1 Introduction

15 Quantum state tomography (QST), the task of estimating an unknown quantum state given
16 measurement outcomes, is essential to building reliable quantum computing devices [52]. The states
17 of the quantum bits (qubits) prepared by an experimental apparatus are estimated, in order to check
18 the correctness of the apparatus and, if needed, determine how to calibrate it. Moreover, quantum
19 process tomography, the task of estimating an unknown quantum channel, can also be cast as a QST
20 problem [7]. There are various approaches to QST, such as trace regression [49, 28, 41, 25, 63, 64],
21 maximum-likelihood (ML) estimation [33, 34, 12], Bayesian estimation [10, 11], and recently
22 proposed deep learning-based methods [3, 53]. Among existing approaches, the ML approach has
23 been widely adopted for its relatively low estimation error in practice and asymptotic statistical
24 guarantees in theory [34, 55].

25 Computing the ML estimator amounts to solving an optimization problem. Whereas the optimization
26 problem is convex, standard convex optimization methods are not directly applicable. It is easily
27 checked that the negative log-likelihood function in ML QST is neither Lipschitz nor smooth,
28 violating standard assumptions in optimization literature [39, 24]. Hence, for example, even whether
29 vanilla gradient descent converges for QST is unclear. **This is perhaps why $R\rho R$, a heuristic
30 algorithm known to be empirically fast, was developed via an expectation maximization, instead
31 of convex optimization, argument [43, 44]. Unfortunately, $R\rho R$ does not always converge [60].** The
32 negative log-likelihood function is indeed self-concordant, so Newton’s method is readily applicable
33 [45]. Nevertheless, the dimension of a quantum state grows exponentially with the number of
34 qubits; the Hessian computations in Newton’s method are computationally too expensive when the
35 dimension is high. There are a few first-order (i.e., gradient-based) convex optimization algorithms

provably converging for ML QST, such as diluted $R\rho R$ [60, 27], SCOPT [57]¹, NoLips [9], the Frank-Wolfe method [48, 24, 65, 14], and entropic mirror descent with line search [39]. These are all batch methods. As they require computing the full gradient in every iteration, their per-iteration time complexities are at least linear in the sample size. To estimate a quantum state, it has been proved that the sample size must be exponential in the number of qubits [47, 30, 17].

Regarding the high dimension and sample size issues in ML QST, it is desirable to, like how we handle the same issues in modern machine learning applications, develop a stochastic first-order optimization method for ML QST. A stochastic first-order optimization method takes one or a few, instead of all, samples in each iteration and avoids computationally expensive Hessian computations. The stochastic quasi-Newton method for self-concordant minimization of Zhou et al. [66] seems to apply. Nevertheless, its step size selection rule involves Hessian computations; moreover, its analysis assumes a bounded Hessian, which does not hold in ML QST. The stochastic mirror-prox and stochastic primal-dual hybrid gradient methods were considered for problems very similar to ML QST [4, 16, 32]. However, their analyses assume either a bounded dual domain or Lipschitzness; both are violated in ML QST.

In this paper, we propose a stochastic first-order algorithm for ML QST. We design the algorithm by an online learning argument. Consider an online convex optimization problem, where the loss function in each round corresponds to the negative log-likelihood function corresponding to one data point in ML QST. Interestingly, this online convex optimization problem is exactly the quantum analogue of online portfolio selection, a celebrated online learning problem [19, 20]. Since the ML approach aims to minimize the empirical average of the negative log-likelihood, once we “quantumize” any existing first-order online portfolio selection algorithm that is no-regret and apply an online-to-batch conversion [15, 22], the resulting algorithm becomes a stochastic first-order algorithm for ML QST. We refer the reader to Section 2 for an introduction of relevant concepts.

The algorithm we choose to “quantumize” is Soft-Bayes [51]. There are two reasons. First, the per-round time complexity of Soft-Bayes is linear in the ambient dimension, arguably the lowest one can expect; second, Soft-Bayes has a curious connection with expectation maximization [43, 44] (see Section 5). We call the resulting algorithm Stochastic Q-Soft-Bayes. Stochastic Q-Soft-Bayes processes one randomly chosen data point in each iteration. Suppose the quantum state to be estimated is represented by a D -by- D density matrix. The per-iteration time complexity of Stochastic Q-Soft-Bayes is $O(D^3)$, independent of the sample size. The expected optimization error of Stochastic Q-Soft-Bayes converges to zero at a $O(\sqrt{(1/T)D \log D})$ rate, where T denotes the number of iterations.

The main technical difficulty lies in figuring out an appropriate quantum extension of Soft-Bayes that coincides with Soft-Bayes when all matrices involved share the same eigenspace and allows for a regret analysis. This is challenging because for any given “non-quantum” expression, one can immediately find many candidates for its quantum extension, but only a few or one of them inherit the desired theoretical properties of their “non-quantum” counterpart; see, e.g., the discussion in [62, Chapter 11] for extending information theoretic quantities to the quantum case. Similar to the quantum extension of exponentiated gradient update by Tsuda et al. [58], the quantum extension we find reveals the complicated mathematical structure of Soft-Bayes hidden in the “non-quantum” setup.

Instead of empirically beating state of the arts, our aim is to give the first provably fast stochastic first-order algorithm for ML QST. Section 3.3 shows that Stochastic Q-Soft-Bayes is competitive in time complexity in comparison to existing batch algorithms. Section 4 shows that Stochastic-Soft-Bayes is empirically even faster than $R\rho R$ in terms of the number of epochs. Unfortunately, Section A shows that in terms of the elapsed time, Q-Soft-Bayes may not be satisfactory to practitioners. We discuss the possibility of developing faster stochastic first-order methods in Section 5.

1.1 Related work

A textbook approach to quantum state tomography is to approximate the problem as a trace regression problem [46] and compute the corresponding least-squares estimate or directly minimize the expected square loss, sometimes with regularization [49, 28, 25, 64, 63]. Since minimizing the square loss is arguably the most standard problem in optimization and machine learning, many existing algorithms apply. Youssry et al. [64] proved the convergence of stochastic entropic mirror descent. Yang et al.

¹A similar algorithm is proposed and studied in [26], but the bounded Hessian assumption therein renders the algorithm inapplicable to ML QST.

[63] showed that several standard online learning algorithms are no-regret for the corresponding online trace regression problem. Notice that both papers do not consider the ML formulation.

Quantum state tomography schemes optimal or nearly optimal in sample complexity are known [47, 30, 38, 29]. The optimal schemes require entangled measurements, challenging to implement [47, 30]. If only incoherent measurements (as in the ML QST scheme considered in this paper) are allowed, the scheme by Kueng et al. [38] is optimal [17]; nevertheless, the scheme is still challenging to implement [38, p. 97]. The scheme proposed by Guță et al. [29] is nearly optimal, but the numerical result in [29, Figure 1] shows that the ML approach achieves a smaller estimation error empirically.

A problem closely related to quantum state tomography is shadow tomography, in which one is not interested in recovering the quantum state but estimating the probability distributions of its measurement outcomes [1]. Aaronson et al. showed that shadow tomography can be done in an online fashion, via follow the regularized leader with the von Neumann entropy [2]. We emphasize that shadow tomography is fundamentally different from quantum state tomography. Indeed, Aaronson showed shadow tomography is strictly easier than state tomography, in the sense that the former requires much less samples than the latter [1]. Another closely related problem is the quantum version of individual sequence prediction considered by Koolen et al. [37]. The loss function studied in [37] is the trace-log loss, instead of the log-trace loss we consider, as discussed in Section 4 of their paper.

Our algorithm is developed via “quantumizing” an online portfolio selection algorithm. Online portfolio selection is a classic online learning problem. It is known that the optimal regret of online portfolio selection is $O(D \log T)$, where D denotes the ambient dimension and T denotes the number of rounds, and is achieved by Universal Portfolio Selection (UPS) [19, 20]. However, UPS is computationally too expensive to be practical [35]. There are several algorithms that try to balance between the regret and computational complexity, but none of them is optimal in both aspects [42, 59]. Soft-Bayes strikes a balance with a $O(D)$ per-round time complexity and $O(\sqrt{TD \log D})$ regret.

Recently, Zimmert et al. [67] “quantumized” another online portfolio selection algorithm, called BISOONS, to solve the game of online quantum state tomography described in Section 3.1². By an online-to-batch conversion, their algorithm yields a stochastic algorithm for ML QST. The resulting algorithm achieves a better iteration complexity than Stochastic Q-Soft-Bayes; nevertheless, each iteration of it requires solving a self-concordant convex program by, e.g., Newton’s method, resulting in a high time complexity incomparable to that of Stochastic Q-Soft-Bayes. In the words of Zimmert et al. [67], both their and our algorithms are on the state-of-the-art efficiency-regret frontier.

1.2 Notations

We write \mathbb{R}_+ for the set of non-negative real numbers and \mathbb{R}_{++} the set of strictly positive real numbers. Let $J \in \mathbb{N}$. We write $[J]$ for the set $\{1, \dots, J\}$. Let M be a matrix. We write M^H for its Hermitian (conjugate transpose) and $\text{tr}(M)$ for its trace. Let H be a Hermitian matrix; we write its spectral decomposition as $H = \sum_d \lambda_d P_d$, where λ_d are the eigenvalues and P_d are projections onto the associated eigenspaces. Let f be a real-valued function whose domain contains $\{\lambda_d\}$. Then, $f(H)$ is defined as the matrix $\sum_d f(\lambda_d) P_d$. Let A and B be two matrices. We write $A \geq B$ if and only if $A - B$ is positive semi-definite. Let \mathcal{E} be an event and ξ be a random variable following a probability distribution P . We write $P(\mathcal{E})$ for the probability of the event and $E_P[\xi]$ for the expectation of ξ . We sometimes omit the subscript P and write $E[\xi]$ when there is no ambiguity.

2 Preliminaries

2.1 Maximum-Likelihood Quantum State Tomography

In the mathematical formulation of quantum mechanics, a quantum state corresponds to a *density matrix*, a Hermitian positive semi-definite complex matrix of unit trace. Let the dimension of the density matrix be $D \in \mathbb{N}$. If there are q qubits, then $D = 2^q$. We denote by \mathcal{D} the set of density matrices in $\mathbb{C}^{D \times D}$, i.e.,

$$\mathcal{D} := \{ \rho \mid \rho \in \mathbb{C}^{D \times D}, \rho = \rho^H, \rho \geq 0, \text{tr } \rho = 1 \}.$$

²We note that the work of Zimmert et al. [67] appears much later than the arXiv version of our work. This footnote is simply to address potential confusions of reviewers and may be removed in the camera-ready version. We do not encourage the reviewers to check the arXiv version as that violates the double-blind policy.

A measurement setup corresponds to a *positive operator-valued measure (POVM)*, a set of Hermitian positive semi-definite complex matrices summing to the identity. Let $\rho \in \mathcal{D}$ and $\{M_1, \dots, M_J\} \subset \mathbb{C}^{D \times D}$ be a POVM. The measurement outcome is a random variable η taking values in $[J]$ such that

$$P(\eta = j) = \text{tr}(M_j \rho), \quad \forall j \in [J].$$

The ML estimation approach seeks the quantum state that maximizes the probability of observing the measured data. Let $\rho^\dagger \in \mathcal{D}$ be the density matrix to be estimated. In a standard QST experiment, we construct N independent copies of ρ^\dagger and measure the copies independently with possibly different POVMs. It is easily checked that the ML estimator is of the form

$$\hat{\rho} \in \underset{\rho \in \mathcal{D}}{\text{argmax}} \prod_{n=1}^N \text{tr}(A_n \rho),$$

where each A_n is an element of the POVM for the n -th measurement. We call $\{A_1, \dots, A_N\}$ the *data-set*. We equivalently write the ML estimator as

$$\hat{\rho} \in \underset{\rho \in \mathcal{D}}{\text{argmin}} f(\rho), \quad (1)$$

$$f(\rho) := \frac{1}{N} \sum_{n=1}^N (-\log \text{tr}(A_n \rho)). \quad (2)$$

Obviously, (1) is a convex optimization problem. If the matrix $A_{n'}$ is not full-rank for some $n' \in [N]$ (as in the cases with the Pauli measurement [40] and Pauli basis measurement [54, 56]), then $\text{tr}(A_{n'} \rho)$ can be arbitrarily close to zero on \mathcal{D} and hence the k -th-order derivative of the objective function f is unbounded for all $k \in \mathbb{N}$.

Let A be a random matrix following the empirical distribution \hat{P}_N on the data-set $\{A_1, \dots, A_N\}$. If the matrices A_n are all different, then \hat{P}_N is simply the uniform distribution on the data-set $\{A_1, \dots, A_N\}$. Then, we can write the objective function in (1) as an expectation

$$f(\rho) = \mathbb{E}_{\hat{P}_N} [-\log \text{tr}(A \rho)]. \quad (3)$$

This observation connects ML QST with the problem of computing the log-optimal portfolio.

2.2 Log-optimal Portfolio

Interestingly, the optimization problem (1) is exactly a quantum extension of the problem of computing the log-optimal portfolio (aka the Kelly criterion), an asymptotically optimal strategy for long-term investment [5, 13, 36]. Consider multi-round investment in a market. Suppose there are D investment alternatives. For the t -th round, we list the return rates of the investment alternatives in that round as a random vector $a_t \in \mathbb{R}_+^D$. Before each round starts, an investor needs to determine the portfolio for the round given the past return rates. Denote by P_{t+1} the probability distribution of a_{t+1} conditional on the history (a_1, \dots, a_t) . The log-optimal portfolio w_{t+1}^* for the $(t+1)$ -th round is given by the stochastic optimization problem:

$$w_{t+1}^* \in \underset{w \in \Delta}{\text{argmin}} \varphi(w), \quad (4)$$

$$\varphi(w) := \mathbb{E}_{P_{t+1}} [-\log \langle a_{t+1}, w \rangle], \quad (5)$$

where Δ denotes the probability simplex in \mathbb{R}^D , the set of entry-wise non-negative vectors whose entries sum to one. Then, the investor distributes the wealth to the investment alternatives following the ratios specified by w_{t+1}^* .

We now discuss the correspondence between ML QST and log-optimal portfolio. The set \mathcal{D} is indeed a quantum extension of the probability simplex Δ , in the sense that a Hermitian matrix is a density matrix if and only if its vector of eigenvalues lies in the probability simplex. The objective functions in (1) and (4) are both expectations of the logarithm of linear functions. Indeed, it is easily checked that if the matrices involved in (1) share the same eigenbasis, then the non-commutativity issue in the quantum setup vanishes and (1) coincides with (4). Though the correspondence is obvious given the two problem formulations, it seems that this correspondence has not been discussed in the literature.

2.3 Online Portfolio Selection

Online portfolio selection may be viewed as a probability-free version of log-optimal portfolio [19]. Online portfolio selection is a multi-round game between two players, say INVESTOR and MARKET. Suppose the game consists of T rounds. In the t -th round of the game, first, INVESTOR announces a portfolio $w_t \in \Delta$; then, MARKET announces the return rates of all investment alternatives for the t -th round in a vector $a_t \in \mathbb{R}_+^D$; finally, INVESTOR suffers a loss of value $-\log \langle a_t, w_t \rangle$. The goal of INVESTOR is to achieve a low *regret* against all possible strategies of MARKET. The regret is defined as

$$R_T := \sup \sum_{t=1}^T (-\log \langle a_t, w_t \rangle) - \min_{w \in \Delta} \sum_{t=1}^T (-\log \langle a_t, w \rangle),$$

where the supremum is over all possible strategies of MARKET to determine $(a_t)_{1 \leq t \leq T}$. We say an algorithm for INVESTOR to determine the portfolios is *no-regret* if it achieves $R_T = o(T)$.

If we can sample from the conditional probability distribution P_{t+1} specified in the previous sub-section, then we can transform a no-regret online portfolio selection algorithm to an algorithm that approximately computes the log-optimal portfolio. The following is an immediate consequence of the online-to-batch conversion [15, 50].

Proposition 2.1. *Suppose in the online portfolio selection game, the vectors a_t are all independent and identically distributed (i.i.d.) random vectors following the probability distribution P_{t+1} in the previous sub-section. Let $(w_t)_{t \in \mathbb{N}}$ be the sequence of iterates generated by an algorithm for INVESTOR of regret R_T . Then, for any $T \in \mathbb{N}$,*

$$\mathbb{E} \left[\varphi(\bar{w}_T) - \min_{w \in \Delta} \varphi(w) \right] \leq \frac{R_T}{T}, \quad \bar{w}_T := \frac{w_1 + \dots + w_T}{T}.$$

Recall that φ is the conditional expectation of the log-linear loss in (5).

If INVESTOR adopts a no-regret algorithm, then the expected optimization error vanishes as $T \rightarrow \infty$.

2.4 Soft-Bayes

There are various existing algorithms for online portfolio selection. Among these algorithms, we are particularly interested in Soft-Bayes [51]. The per-iteration time complexity of Soft-Bayes is linear in D , arguably the lowest one can expect. This is a desirable feature for ML QST, as the dimension of the density matrix grows exponentially with the number of qubits.

The iteration rule of Soft-Bayes is as follows.

- Initialize at $w_1 = (1/D, \dots, 1/D) \in \Delta$ (the uniform distribution).
- For each $t \in \mathbb{N}$, compute

$$w_{t+1} = (1 - \eta)w_t + \eta \frac{a_t \circ w_t}{\langle a_t, w_t \rangle}, \quad \forall t \in \mathbb{N}, \quad (6)$$

for some properly chosen *learning rate* $\eta \in [0, 1]$, where \circ denotes the entry-wise product.

Soft-Bayes has the following regret guarantee.

Theorem 2.2 ([51]). *After T rounds in online portfolio selection, the regret of Soft-Bayes with*

$$\eta = \frac{\sqrt{DT}}{\sqrt{DT} + \sqrt{\log D}} \quad (7)$$

is at most $2\sqrt{TD \log D} + \log D$.

3 Online Maximum-Likelihood Quantum State Tomography by Q-Soft-Bayes

Following the discussion in Section 1, we first propose a *game of online quantum state tomography* as a quantum extension of online portfolio selection. Then, we “quantumize” Soft-Bayes to derive a no-regret algorithm for the game and analyse its regret. Finally, we adopt the online-to-batch conversion and bound the expected optimization error of the resulting algorithm.

3.1 Game of Online Quantum State Tomography

We propose the following game of online quantum state tomography as a quantum extension of online portfolio selection. Online quantum state tomography is a multi-round game between two players, say PHYSICIST and ENVIRONMENT. Suppose there are in total T rounds. In the t -th round, first, PHYSICIST announces a density matrix $\rho_t \in \mathcal{D}$; then, ENVIRONMENT announces a Hermitian positive semi-definite matrix $A_t \geq 0$; finally, PHYSICIST suffers for a loss of value $-\log \text{tr}(A_t \rho_t)$. The regret in this game is given by

$$R_T := \sup \sum_{t=1}^T (-\log \text{tr}(A_t \rho_t)) - \min_{\gamma \in \mathcal{D}} \sum_{t=1}^T (-\log \text{tr}(A_t \gamma)),$$

where the supremum is over all possible strategies of PHYSICIST to generate the sequence $(A_t)_{1 \leq t \leq T}$.

The connection with online portfolio selection is obvious and similar to that between ML QST and the log-optimal portfolio: The vector of eigenvalues of a density matrix lies in the probability simplex Δ ; the Hermitian matrices A_t and the positive semi-definiteness condition correspond to the vectors a_t in online portfolio selection and their non-negativity condition, respectively; the losses in the two games are both logarithms of linear functions; the regrets in the two games are defined exactly in the same manner. When all the matrices involved in the game of online quantum state tomography share the same eigenbasis, we recover the game of online portfolio selection.

3.2 Q-Soft-Bayes and Stochastic Q-Soft-Bayes

We propose the following Q-Soft-Bayes algorithm as a quantum extension of Soft-Bayes.

- Initialize at $\rho_1 = W_1 = I/D$.
- For each $t \in \mathbb{N}$, compute

$$\begin{aligned} G_t &= (1 - \eta)I + \eta \frac{A_t}{\text{tr}(A_t \rho_t)}, \\ W_{t+1} &= \exp(\log(W_t) + \log(G_t)), \\ \rho_{t+1} &= \frac{W_{t+1}}{\text{tr}(W_{t+1})}, \end{aligned} \tag{8}$$

for some properly chosen learning rate $\eta \in [0, 1]$.

Remark 3.1. Recently, we learned that Q-Soft-Bayes may be interpreted using the *commutative matrix product* by Warmuth and Kuzmin [61]. It is currently unclear to us whether this interpretation provides any insight.

If we were able to cancel the exponential and logarithms in Q-Soft-Bayes, then we recover Soft-Bayes; however, due to the non-commutativity issue, such cancellation is illegal in general. In comparison to Soft-Bayes, Q-Soft-Bayes has an additional normalization step to ensure its outputs are of unit trace. We prove the following in Section B.

Proposition 3.2. *It holds that $\text{tr}(W_t) \leq 1$ for all t .*

Numerical experiments show that the equality does not always hold, so the normalization step is necessary. Recall that Soft-Bayes does not need the normalization step (see Section 2.4).

In Appendix C, we prove the following regret bound for Q-Soft-Bayes, showing that it inherits the regret bound of Soft-Bayes.

Theorem 3.3. *The regret of Q-Soft-Bayes with the learning rate η given in (7) is at most $2\sqrt{TD} \log D + \log D$.*

Remark 3.4. One might wonder why D is not replaced by D^2 in the quantum case. This is because in our extension, the analogue of a D -dimensional vector in the quantum case is the D -dimensional vector of eigenvalues of a D -by- D Hermitian matrix, instead of the D -dimensional vector obtained by vectorizing a \sqrt{D} -by- \sqrt{D} matrix. Similar coincidence of regret bounds can be observed in, for example, the matrix version of exponentiated gradient update [58, 8] and the quantum individual sequence prediction algorithms of Koolen et al. [37].

247 The standard online-to-batch conversion argument can also be applied to solving ML QST by
 248 Q-Soft-Bayes. Recall for ML QST, our aim is to solve the stochastic optimization problem

$$\hat{\rho} \in \operatorname{argmin}_{\rho \in \mathcal{D}} \mathbb{E}_{\hat{P}_N} [-\log \operatorname{tr}(A, \rho)],$$

249 where A is a random matrix following the empirical probability distribution \hat{P}_N on the data-set
 250 $\{A_1, \dots, A_N\}$ (see Section 2.1). We propose the following stochastic optimization algorithm, which
 251 we call Stochastic Q-Soft-Bayes, to solve ML QST.

- 252 • Initialize Q-Soft-Bayes with $\rho_1 = W_1 = I/D$.
- 253 • In the t -th iteration of Stochastic Q-Soft-Bayes, do the following.
 - 254 1. Output the t -th output ρ_t of Q-Soft-Bayes.
 - 255 2. Sample a random matrix $B_t \in \{A_1, \dots, A_N\}$ following the empirical probability
 256 distribution \hat{P}_N on the data-set, independent of the past.
 - 257 3. Let ENVIRONMENT in the online QST game announce the matrix B_t .

258 Similarly as for Proposition 2.1, the standard online-to-batch conversion argument provides the
 259 following convergence guarantee of Stochastic-Q-Soft-Bayes.

260 **Proposition 3.5.** *Let $(\rho_t)_{t \in \mathbb{N}}$ be the sequence of iterates generated by Stochastic Soft-Bayes. Then,*
 261 *for any $T \in \mathbb{N}$, it holds that*

$$\mathbb{E} \left[f(\bar{\rho}_T) - \min_{\rho \in \mathcal{D}} f(\rho) \right] \leq 2\sqrt{\frac{D \log D}{T}} + \frac{\log D}{T},$$

262 where $\bar{\rho}_T := (\rho_1 + \dots + \rho_T)/T$ and the expectation is with respect to the randomness in B_t of
 263 Stochastic Soft-Bayes. Recall f is the objective function in ML QST as defined in (2) or (3) (the two
 264 definitions are equivalent).

265 Therefore, Stochastic-Q-Soft-Bayes outputs an approximate ML estimator of expected optimization
 266 error smaller than ε in $O((D \log D)/\varepsilon^2)$ iterations. Each iteration of Stochastic-Q-Soft-Bayes
 267 requires computing a matrix exponential and two matrix logarithms. The overall time complexity is
 268 hence $O((D^4 \log D)/\varepsilon^2)$. One may adopt anytime online-to-batch [22], which seems to empirically
 269 yield faster convergence. According to [22], the optimization error guarantee remains the same; the
 270 only difference is that ∇f are evaluated at $\bar{\rho}_t$ instead of ρ_t when implementing Soft-Bayes, so the
 271 overall time complexity also remains the same.

272 One may be interested in the distance to the minimizer. It is easily checked that the function f is
 273 self-concordant. If $\nabla^2 f$ is positive definite at the minimizer, a standard condition for well-posed
 274 estimators, then the function f is locally strongly convex around the minimizer [45, Theorem 4.1.6].
 275 Therefore, the distance to the minimizer, measured in terms of the Frobenius norm, is asymptotically
 276 of the order of the square root of the optimization error.

277 3.3 Theoretical Comparison with Existing Batch Algorithms

278 Let us compare the time complexities of Stochastic Q-Soft-Bayes and existing algorithms discussed
 279 in Section 1. The iteration complexities of existing algorithms are mostly unknown or vague in
 280 their dependence on the problem parameters. Diluted $R\rho R$ and entropic mirror descent with line
 281 search do not have non-asymptotic analysis results [60, 27, 39]; SCOPT only has a local linear
 282 rate guarantee [57]; Adaptive Frank-Wolfe and Monotonous Frank-Wolfe have $O(\varepsilon^{-1})$ iteration
 283 complexities with unclear dependence on the dimension and sample size, as their error bounds involve
 284 local smoothness parameters that are hard to evaluate [14, 24]. A finer analysis of Adaptive Frank-
 285 Wolfe by Zhao and Freund [65] shows that its iteration complexity is $O(\varepsilon^{-1}N)$ and hence its time
 286 complexity is $O(\varepsilon^{-1}(N^2 D^2 + N\tau))$, where the symbol τ denotes the time of computing the local
 287 norm defined by the Hessian, for which we do not know an efficient implementation. In comparison,
 288 the complexities of Stochastic Q-Soft-Bayes is very clear: $O(\varepsilon^{-2} D \log D)$ iteration complexity and
 289 hence $O(\varepsilon^{-2} D^4 \log D)$ time complexity. The time complexity of Stochastic Q-Soft-Bayes becomes
 290 competitive with Adaptive Frank-Wolfe if $N \gg D\sqrt{(1/\varepsilon) \log D}$, ignoring the time of computing the
 291 local norms. Recently, it is proved that any QST scheme with non-coherent measurement, e.g., ML
 292 QST we consider in this paper, requires $N = \Omega(D^3/\delta^2)$ to achieve an estimation error smaller than δ

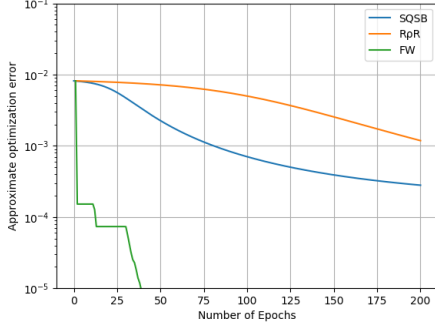


Figure 1: Approximate optimization errors in function value of Stochastic Q-Soft-Bayes (SQSB), $R\rho R$ (~~RrhoR~~), and Monotonous Frank-Wolfe (FW).

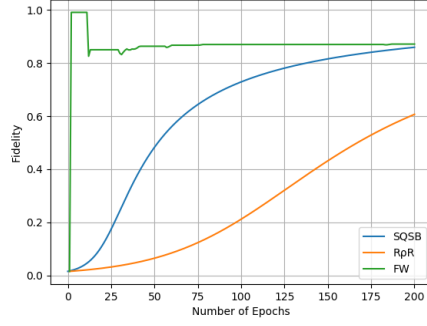


Figure 2: Fidelity values of the iterates and the W state achieved by Stochastic Q-Soft-Bayes (SQSB), $R\rho R$ (~~RrhoR~~), and Monotonous Frank-Wolfe (FW).

in the trace distance [17]. The algorithm by Zimmert et al. [67] has a $\tilde{O}(D^3/\varepsilon)$ iteration complexity and $O(D^6)$ per-iteration time complexity ignoring the dependence on other parameters, due to the use of Newton’s method to compute the iterates; the overall time complexity has a much higher dependence on the dimension than Adaptive Frank-Wolfe and Stochastic Q-Soft-Bayes. We conclude that the time complexity of Stochastic Q-Soft-Bayes is competitive compared to existing algorithms.

4 Numerical Results

As discussed above, Stochastic Q-Soft-Bayes is competitive in theory. We now examine its empirical performance with anytime online-to-batch. We compare its empirical speed with two batch methods, the $R\rho R$ method [43, 44] and Monotonous Frank-Wolfe [14], on a synthetic data-set in Figure 1 and Figure 2. We have mentioned several batch methods applicable for ML QST in Section 1. Among them, we choose $R\rho R$ for comparison as it is representative in physics literature and empirically fast, though it does not always converge. We choose monotonous Frank-Wolfe for comparison as it avoids computationally expensive Hessian computations in step size selection. Recall that Monotonous Frank-Wolfe converges at a $O(1/t)$ rate as other Frank-Wolfe methods for self-concordant minimization do [65, 24, 48], but its complexity guarantee lacks a clear characterization of the dependence on the dimension and sample size.

The synthetic data-set is generated basically following the set-up in [31]. The number of qubits q equals 6. The dimension D then equals $2^q = 64$. The unknown quantum state to be measured is the W state. We randomly generate $N = 4^q \times 100 = 409600$ Pauli observables as in, e.g., [25, 28, 40], each of which corresponds to a POVM of two rank- $(D/2)$ elements. As there are in total 4^q different Pauli observables (and hence POVMs), each observable is used about 100 times. Then, we sample the N measurement outcomes and formulate the ML estimator following Section 2.1.

The performance measures we consider are optimization errors (in objective function) and fidelity values. To estimate the optimization error, we run each algorithm for 200 epochs and use the smallest function value found by the algorithms as an approximate optimal value. The approximate optimization error of an iterate is defined as the difference between the objective function value at the iterate and the approximate optimal value. Fidelity is a notion commonly used by physicists to measure how close two quantum states are to each other. For any two density matrices ρ and σ , the fidelity is given by $F(\rho, \sigma) := (\text{tr} \sqrt{\sqrt{\rho}\sigma\sqrt{\rho}})^2$, which takes values in $[0, 1]$. The fidelity of two quantum states equals 1, if the two states are exactly the same. We plot the optimization errors and fidelity values versus the number of epochs. An epoch corresponds to one pass of the whole data-set. One iteration of Stochastic Q-Soft-Bayes corresponds to $1/N$ epoch. One iteration of $R\rho R$ and Monotonous Frank-Wolfe corresponds to 1 epoch as both algorithms are batch.

Obviously, Stochastic Q-Soft-Bayes converges faster than $R\rho R$ in both optimization error and fidelity. Where as Monotonous Frank-Wolfe is the fastest in both figures, this can be explained by the fact that

Frank-Wolfe tends to generate approximately low-rank iterates. The W state corresponds to a rank-1 density matrix, so the ML estimate should be approximately low-rank, matching the structure favoured by Frank-Wolfe. We conclude that the convergence speed of Stochastic Q-Soft-Bayes is competitive in theory (Section 3.3) and comparable to fast yet theoretically non-rigorous algorithms in practice.

A comparison in terms of the elapsed time is provided in Appendix A. The results show there is a large room for improvement to compete with $R\rho R$ and Monotonous Frank-Wolfe in the elapsed time. The source codes are provided in the supplementary material.

5 Discussions

5.1 Can We Find a Faster Stochastic First-Order Algorithms for ML QST?

Our approach to constructing a stochastic first-order algorithm for ML QST conceptually applies to any no-regret online portfolio selection algorithm. In this paper, we focus on Soft-Bayes. Other existing online portfolio selection algorithms have much higher per-iteration time complexities, in terms of the dependence on the ambient dimension and sample size. If we adopt any other existing online portfolio selection algorithm and “quantumize” it to obtain a stochastic algorithm for ML QST, then the resulting algorithm will scale poorly with the number of qubits. Developing an online portfolio selection algorithm that enjoys both a low regret and low time complexity is still open [59, 67].

It is still possible to develop another quantum extension of Soft-Bayes that enjoys a lower per-iteration time complexity. The per-iteration time complexity issue may be mitigated if we consider other quantum extensions of Soft-Bayes. For example, if we naïvely replace (8) by $W_{t+1} = (G_t W_t + W_t G_t) / 2$, the resulting algorithm still coincides with Soft-Bayes when all matrices share the same eigenbasis, whereas the per-iteration time complexity is reduced to $O(D^\omega)$ for some $\omega < 2.373$ [6]. Unfortunately, we cannot work out a non-asymptotic analysis for any other possible quantum extension of Soft-Bayes we can think of.

The discussion above assumes that we adopt the online-to-batch argument as in this paper. Another way, which we think perhaps more plausible, is to directly consider the stochastic optimization formulation and develop a stochastic optimization algorithm for ML QST.

5.2 Connection with Expectation Maximization

Finally, let us discuss an interesting connection between Q-Soft-Bayes and expectation maximization (EM). The $R\rho R$ algorithm, according to [43, 44], was inspired by the expectation maximization (EM) method for solving optimization problems of the form (4). Given a full-rank initial iterate $\rho_1 \in \mathcal{D}$, $R\rho R$ iterates as

$$\rho_{t+1} = \frac{R_t \rho_t R_t}{\text{tr}(R_t \rho_t R_t)}, \quad R_t := -\nabla f(\rho_t), \quad \forall t \in \mathbb{N},$$

where f is defined in (2). In comparison, given an entry-wise positive vector $w_1 \in \Delta$, EM for (4) iterates as

$$w_{t+1} = w_t \circ (-\nabla \varphi(w_t)), \quad \forall t \in \mathbb{N}.$$

It is interesting to notice that even when all matrices involved share the same eigenbasis, $R\rho R$ is not equivalent to EM. Indeed, EM is proved to asymptotically converge to the optimum [18, 21], whereas $R\rho R$ oscillates on a carefully designed data-set [60]. This suggests that $R\rho R$ is perhaps not a “natural” quantum extension of EM. Later, there were variations of $R\rho R$ that solve the convergence issue by line search [60, 27], but these variations still do not recover EM.

Notice that the formulation of Soft-Bayes (6) is the convex combination of the previous iterate and the output of EM. Therefore, Soft-Bayes, after the online-to-batch conversion, can be interpreted as a relaxed stochastic EM method for computing the log-optimal portfolio. As Q-Soft-Bayes becomes Soft-Bayes when all matrices involved share the same eigenbasis, we may claim that Stochastic Q-Soft-Bayes is also a relaxed stochastic EM method, though its derivation does not have any obvious relation with the standard derivation of EM [23].

References

- [1] S. Aaronson. Shadow tomography of quantum states. *SIAM J. Comput.*, 49(5):STOC18–368–STOC18–394, 2020.
- [2] S. Aaronson, X. Chen, E. Hazan, S. Kale, and A. Nayak. Online learning of quantum states. In *Adv. Neural Information Processing Systems 31*, 2018.
- [3] S. Ahmed, C. S. Muñoz, F. Nori, and A. F. Kockum. Quantum state tomography with conditional generative adversarial networks, 2020. arXiv:2008.03240.
- [4] A. Alacaoglu. *Adaptation in Stochastic Algorithms: From Nonsmooth Optimization to Min-Max Problems and Beyond*. PhD thesis, École polytechnique fédérale de Lausanne, 2021.
- [5] P. H. Algoet and T. M. Cover. Asymptotic optimality and asymptotic equipartition properties of log-optimum investment. *Ann. Probab.*, 16(2):876–898, 1988.
- [6] J. Alman and V. V. Williams. A refined laser method and faster matrix multiplication. In *Proc. 2021 ACM-SIAM Symp. Discrete Algorithms (SODA)*, 2021.
- [7] J. B. Altepeter, D. Branning, E. Jeffrey, T. C. Wei, P. G. Kwiat, R. T. Thew, J. L. O’Brien, M. A. Nielsen, and A. G. White. Ancilla-assisted quantum process tomography. *Phys. Rev. Lett.*, 90(19), 2003.
- [8] S. Arora, E. Hazan, and S. Kale. The multiplicative weights update method: A meta-algorithm and applications. *Theory Comput.*, 8:121–164, 2012.
- [9] H. H. Bauschke, J. Bolte, and M. Teboulle. A descent lemma beyond Lipschitz gradient continuity: first-order methods revisited and applications. *Math. Oper. Res.*, 42(2):330–348, 2017.
- [10] R. Blume-Kohout. Hedged maximum likelihood quantum state estimation. *Phys. Rev. Lett.*, 105, 2010.
- [11] R. Blume-Kohout. Optimal, reliable estimation of quantum states. *New J. Phys.*, 12, 2010.
- [12] E. Bolduc, G. C. Knee, E. M. Gauger, and J. Leach. Projected gradient descent algorithms for quantum state tomography. *npj Quantum Inf.*, 3, 2017.
- [13] L. Breiman. Investment policies for expanding business optimal in a long-run sense. In W. T. Ziemba and R. G. Vickson, editors, *Stochastic Optimization Models in Finance*, pages 593–598. Academic Press, New York, NY, 1975.
- [14] A. Carderera, M. Besançon, and S. Pokutta. Simple steps are all you need: Frank-Wolfe and generalized self-concordant functions, 2021.
- [15] N. Cesa-Bianchi, A. Conconi, and C. Gentile. On the generalization ability of on-line learning algorithms. *IEEE Trans. Inf. Theory*, 50(9):2050–2057, 2004.
- [16] A. Chambolle, M. J. Ehrhardt, P. Richtárik, and C.-B. Schönlieb. Stochastic primal-dual hybrid gradient algorithm with arbitrary sampling and imaging applications. *SIAM J. Optim.*, 28(4):2783–2808, 2018.
- [17] S. Chen, B. Huang, J. Li, A. Liu, and M. Selke. Tight bounds for state tomography with incoherent measurements. 2022. arXiv:2206.05265v1.
- [18] T. M. Cover. An algorithm for maximizing expected log investment return. *IEEE Trans. Inf. Theory*, IT-30(2):369–373, 1984.
- [19] T. M. Cover. Universal portfolios. *Math. Financ.*, 1(1):1–29, 1991.
- [20] T. M. Cover and E. Ordentlich. Universal portfolios with side information. *IEEE Trans. Inf. Theory*, 42(2):348–363, 1996.
- [21] I. Csiszár and G. Tusnády. Information geometry and alternating minimization procedures. *Stat. Decis.*, (Supplement 1):205–237, 1984.

- [22] A. Cutkosky. Anytime online-to-batch, optimism and acceleration. In *Proc. 36th Int. Conf. Machine Learning*, 2019.
- [23] A. P. Dempster, N. M. Laird, and D. B. Rubin. Maximum likelihood from incomplete data via the EM algorithm. *J. R. Stat. Soc., Ser. B*, 39(1):1–38, 1977.
- [24] P. Dvurechensky, P. Ostroukhov, K. Safin, S. Shtern, and M. Staudigl. Self-concordant analysis of Frank-Wolfe algorithms. In *Proc. 37th Int. Conf. Machine Learning*, 2020.
- [25] S. T. Flammia, D. Gross, Y.-K. Liu, and J. Eisert. Quantum tomography via compressed sensing: Error bounds, sample complexity and efficient estimators. *New J. Phys.*, 14, 2012.
- [26] W. Gao and D. Goldfarb. Quasi-Newton methods: superlinear convergence without line searches for self-concordant functions. *Optim. Methods Softw.*, 34(1):194–217, 2019.
- [27] D. S. Gonçalves, M. A. Gomes-Ruggiero, and C. Lavor. Global convergence of diluted iterations in maximum-likelihood quantum tomography. *Quantum Inf. Comput.*, 14(11&12):966–980, 2014.
- [28] D. Gross, Y.-K. Liu, S. T. Flammia, S. Becker, and J. Eisert. Quantum state tomography via compressed sensing. *Phys. Rev. Lett.*, 105, 2010.
- [29] M. Guță, J. Kahn, R. Kueng, and J. A. Tropp. Fast state tomography with optimal error bounds. *J. Phys. A: Math. Theor.*, 53, 2020.
- [30] J. Haah, A. W. Harrow, Z. Ji, X. Wu, and N. Yu. Sample-optimal tomography of quantum states. *IEEE Trans. Inf. Theory*, 63(9):5628–5641, 2017.
- [31] H. Häffner, W. Hänsel, C. F. Roos, J. Benhelm, D. Cheek-al-kar, M. Chwalla, T. Körber, U. D. Rapol, M. Riebe, P. O. Schmidt, C. Becher, O. Gühne, W. Dür, and R. Blatt. Scalable multiparticle entanglement of trapped ions. *Nature*, 438:643–646, 2005.
- [32] N. He, Z. Harchaoui, Y. Wang, and L. Song. Point process estimation with mirror prox algorithms. *Appl. Math. Optim.*, 2019.
- [33] Z. Hradil. Quantum-state estimation. *Phys. Rev. A*, 55(3), 1997.
- [34] Z. Hradil, J. Řeháček, J. Fiurášek, and M. Ježek. Maximum-likelihood methods in quantum mechanics. In *Quantum State Estimation*, chapter 3, pages 59–112. Springer, Berlin, 2004.
- [35] A. Kalai and S. Vempala. Efficient algorithms for universal portfolios. *J. Mach. Learn. Res.*, 3: 423–440, 2002.
- [36] J. L. Kelly, Jr. A new interpretation of information rate. *IRE Trans. Inf. Theory*, 2(3):185–189, 1956.
- [37] W. M. Koolen, W. Kotłowski, and M. K. Warmuth. Learning eigenvectors for free. In *Adv. Neural Information Processing Systems 24*, 2011.
- [38] R. Kueng, H. Rauhut, and U. Terstiege. Low rank matrix recovery from rank one measurements. *Appl. Comput. Harmon. Anal.*, 42:88–116, 2017.
- [39] Y.-H. Li and V. Cevher. Convergence of the exponentiated gradient method with Armijo line search. *J. Optim. Theory Appl.*, 181(2):588–607, 2019.
- [40] Y.-K. Liu. Universal low-rank matrix recovery from Pauli measurements. In *Adv. Neural Information Processing Systems 24*, 2011.
- [41] Y.-K. Liu. Universal low-rank matrix recovery from Pauli measurements. 2011. arXiv:1103.2816v2 [quant-ph].
- [42] H. Luo, C.-Y. Wei, and K. Zheng. Efficient online portfolio with logarithmic regret. In *Adv. Neural Information Processing Systems 31*, 2018.
- [43] A. I. Lvovsky. Iterative maximum-likelihood reconstruction in quantum homodyne tomography. *J. Opt. B: Quantum Semiclass. Opt.*, 6, 2004.

- [44] G. Molina-Terriza, A. Vaziri, J. Řeháček, Z. Hradil, and A. Zeilinger. Triggered qutrits for quantum communication protocols. *Phys. Rev. Lett.*, 92(16), 2004.
- [45] Y. Nesterov. *Introductory Lectures on Convex Optimization*. Kluwer, Boston, MA, 2004.
- [46] M. A. Nielsen and I. L. Chuang. *Quantum Computation and Quantum Information*. Cambridge Univ. Press, Cambridge, UK, 2010.
- [47] R. O’Donnell and J. Wright. Efficient quantum tomography. In *Proc. 48th Annu. ACM Symp. Theory of Computing*, pages 899–912, 2016.
- [48] G. Odor, Y.-H. Li, A. Yurtsever, Y.-P. Hsieh, M. El Halabi, Q. Tran-Dinh, and V. Cevher. Frank-Wolfe works for non-Lipschitz continuous gradient objectives: Scalable Poisson phase retrieval. In *IEEE Int. Conf. Acoustics, Speech and Signal Processing*, pages 6230–6234, 2016.
- [49] T. Opatrný, D.-G. Welsch, and W. Vogel. Least-squares inversion for density-matrix reconstruction. *Phys. Rev. A*, 56(3), 1997.
- [50] F. Orabona. A modern introduction to online learning. 2019. arXiv:1912.13213v1.
- [51] L. Orseau, T. Lattimore, and S. Legg. Soft-Bayes: Prod for mixtures of experts with log-loss. In *Proc. 28th Int. Conf. Algorithmic Learning Theory*, pages 372–399, 2017.
- [52] M. Paris and J. Řeháček, editors. *Quantum State Estimation*. Springer, Berlin, 2004.
- [53] Y. Quek, S. Fort, and H. K. Ng. Adaptive quantum state tomography with neural networks. *npj Quantum Inf.*, 7, 2021.
- [54] C. A. Riofrío, D. Gross, S. T. Flammia, T. Monz, D. Nigg, R. Blatt, and J. Eisert. Experimental quantum compressed sensing for a seven-qubit system. *Nature Commun.*, 2017.
- [55] T. L. Scholten and R. Blume-Kohout. Behavior of maximum likelihood in quantum state tomography. *New J. Phys.*, 20, 2018.
- [56] A. Steffens, C. A. Riofrío, W. McCutcheon, I. Roth, B. A. Bell, A. McMillan, M. S. Tame, J. G. Rarity, and J. Eisert. Experimentally exploring compressed sensing quantum tomography. *Quantum Sci. Tech.*, 2, 2017.
- [57] Q. Tran-Dinh, A. Kyrillidis, and V. Cevher. Composite self-concordant minimization. *J. Mach. Learn. Res.*, 16:371–416, 2015.
- [58] K. Tsuda, G. Rätsch, and M. K. Warmuth. Matrix exponentiated gradient updates for on-line learning and Bregman projection. *J. Mach. Learn. Res.*, 6:995–1018, 2005.
- [59] T. van Erven, D. van der Hoeven, W. Kotłowski, and W. M. Koolen. Open problem: Fast and optimal online portfolio selection. In *Proc. 33rd Conf. Learning Theory*, 2020.
- [60] J. Řeháček, Z. Hradil, E. Knill, and A. I. Lvovsky. Diluted maximum-likelihood algorithm for quantum tomography. *Phys. Rev. A*, 75, 2007.
- [61] M. K. Warmuth and D. Kuzmin. Bayesian generalized probability calculus for density matrices. *Mach. Learn.*, 78:63–101, 2010.
- [62] M. Wilde. From classical to quantum Shannon theory. 2019. arXiv:1106.1445v8.
- [63] F. Yang, J. Jiang, J. Zhang, and X. Sun. Revisiting online quantum state learning. In *Proc. AAAI Conf. Artificial Intelligence*, 2020.
- [64] A. Youssry, C. Ferrie, and M. Tomamichel. Efficient online quantum state estimation using a matrix-exponentiated gradient method. *New J. Phys.*, 21(033006), 2019.
- [65] R. Zhao and R. M. Freund. Analysis of the Frank-Wolfe method for logarithmically-homogeneous barriers, with an extension, 2020.
- [66] C. Zhou, W. Gao, and D. Goldfarb. Stochastic adaptive quasi-Newton methods for minimizing expected values. In *Proc. 34th Int. Conf. Machine Learning*, 2017.
- [67] J. Zimmert, N. Agarwal, and S. Kale. Pushing the efficiency-regret Pareto frontier for online learning of portfolios and quantum states. 2022. arXiv:2202.02765v1.

Checklist

1. For all authors...

- (a) Do the main claims made in the abstract and introduction accurately reflect the paper's contributions and scope? [Yes]
- (b) Did you describe the limitations of your work? [Yes] See Section 4. Whereas the proposed algorithm is fast in theory, its speed is unfortunately not satisfactory to practitioners.
- (c) Did you discuss any potential negative societal impacts of your work? [No] This is a theory work.
- (d) Have you read the ethics review guidelines and ensured that your paper conforms to them? [Yes]

2. If you are including theoretical results...

- (a) Did you state the full set of assumptions of all theoretical results? [Yes] Indeed, our algorithm applies to every data-set generated following Section 2.1.
- (b) Did you include complete proofs of all theoretical results? [Yes] See the appendix.

3. If you ran experiments...

- (a) Did you include the code, data, and instructions needed to reproduce the main experimental results (either in the supplemental material or as a URL)? [Yes] The codes are attached.
- (b) Did you specify all the training details (e.g., data splits, hyperparameters, how they were chosen)? [Yes] See Section 4.
- (c) Did you report error bars (e.g., with respect to the random seed after running experiments multiple times)? [No] The comparison is obvious, so we do not feel it necessary to report error bars.
- (d) Did you include the total amount of compute and the type of resources used (e.g., type of GPUs, internal cluster, or cloud provider)? [Yes] See Section A. This information is not necessary for Section 4.

4. If you are using existing assets (e.g., code, data, models) or curating/releasing new assets...

- (a) If your work uses existing assets, did you cite the creators? [N/A]
- (b) Did you mention the license of the assets? [N/A]
- (c) Did you include any new assets either in the supplemental material or as a URL? [N/A]
- (d) Did you discuss whether and how consent was obtained from people whose data you're using/curating? [N/A]
- (e) Did you discuss whether the data you are using/curating contains personally identifiable information or offensive content? [N/A]

5. If you used crowdsourcing or conducted research with human subjects...

- (a) Did you include the full text of instructions given to participants and screenshots, if applicable? [N/A]
- (b) Did you describe any potential participant risks, with links to Institutional Review Board (IRB) approvals, if applicable? [N/A]
- (c) Did you include the estimated hourly wage paid to participants and the total amount spent on participant compensation? [N/A]

552 A Supplementary Numerical Results

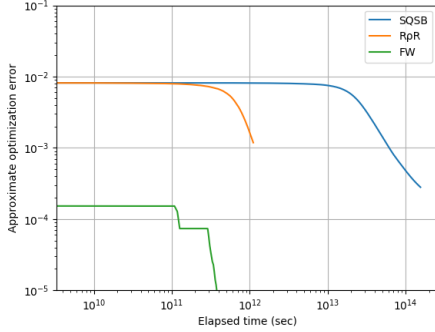


Figure 3: Approximate optimization errors in function value of Stochastic Q-Soft-Bayes (SQSB), $R\rho R$ (**RrhoR**), and Monotonous Frank-Wolfe (FW). Notice that the x-axis is the elapsed time instead of number of epochs.

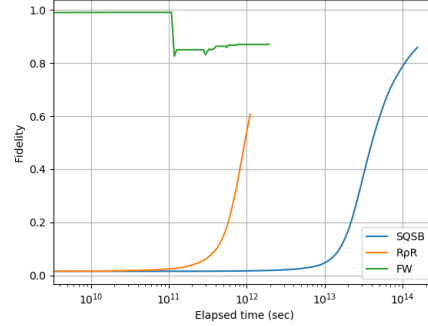


Figure 4: Fidelity values of the iterates and the W state achieved by Stochastic Q-Soft-Bayes (SQSB), $R\rho R$ (**RrhoR**), and Monotonous Frank-Wolfe (FW). Notice the x-axis is the elapsed time instead of number of epochs.

553 As the three algorithms we compare in Section 4 have different per-iteration computational com-
 554 plexities, it can be misleading to evaluate their convergence speeds only in terms of the number of
 555 epochs. Figure 3 and Figure 4 are plotted in terms of the real elapsed time on a [server](#) with a [Intel](#)
 556 [Xeon Gold 5218 CPU of 2.30GHz and 131621512kB memory](#). The algorithms are implemented
 557 with the programming language [Julia](#) and the [Math Kernel Library \(MKL\)](#). The number of threads
 558 for BLAS is set to 8.

559 It is easy to evaluate the per-iteration time complexities. Every iteration of Stochastic Q-Soft-Bayes
 560 involves matrix exponential and matrix logarithm computations, so its per-iteration time complexity
 561 is $O(D^3)$. The time to compute the batch gradient is $O(ND^2)$. Therefore, the per-iteration time
 562 complexity of Monotonous Frank-Wolfe is $O(ND^2)$, dominated by the time of computing the
 563 batch gradient; the per-iteration time complexity of $R\rho R$ is $O(ND^2 + D^{\omega'})$, where the $D^{\omega'}$ term
 564 comes from the time of computing matrix multiplications. Here, one iteration of Stochastic Q-
 565 Soft-Bayes corresponds to $1/N$ “iteration” of $R\rho R$ and Monotonous Frank-Wolfe. However, a
 566 theoretical comparison as in Section 3.3 is impossible. $R\rho R$ does not always converge and hence
 567 lacks an iteration complexity guarantee; the dependence of the iteration complexity of Monotonous
 568 Frank-Wolfe on N and D is unclear.

569 The numerical results show that in terms of the elapsed time, Stochastic Q-Soft-Bayes is much slower
 570 than the other two algorithms, showing a large room for improvement to compete with practical—yet
 571 theoretically non-rigorous—algorithms. Section 5.1 discusses possible approaches to develop faster
 572 algorithms. It should be emphasized again that the elapsed time highly depends on the hardware and
 573 implementation details.

574 B Proof of Proposition 3.2

575 By the Golden-Thompson inequality, we write

$$\begin{aligned} \text{tr}(W_{t+1}) &\leq \text{tr} W_t \left((1 - \eta)I + \eta \frac{A_t}{\text{tr}(A_t \rho_t)} \right) \\ &= (1 - \eta) \text{tr}(W_t) + \eta \frac{\text{tr}(A_t W_t)}{\text{tr}(A_t \rho_t)} \\ &= (1 - \eta) \text{tr}(W_t) + \eta \text{tr}(W_t) \\ &= \text{tr}(W_t), \end{aligned}$$

576 showing that $\text{tr}(W_{t+1}) \leq \text{tr}(W_1) = 1$.

577 C Proof of Theorem 3.3

578 For convenience, we define $\langle A, B \rangle_{\text{H}} = \text{tr}(AB)$ for any two Hermitian matrices A and B . Indeed,
 579 the inner product $\langle \cdot, \cdot \rangle_{\text{H}}$ is the Hilbert-Schmidt inner product in literature. The following lemma—a
 580 reverse Jensen’s inequality for the logarithmic function—is a quantum extension of Lemma 4 in [51].
 581 We omit the proof as it is simply the original proof in [51] applied to the eigenvalues of the matrices
 582 involved.

583 **Lemma C.1.** *Let $\eta \in (0, 1)$, $\rho \in \mathcal{D}$, and $X \in \mathbb{C}^{D \times D}$ be a Hermitian positive semi-definite matrix.*
 584 *We have*

$$\log \langle X, \rho \rangle_{\text{H}} \leq \frac{1}{\eta} \langle \log((1 - \eta)I + \eta X), \rho \rangle_{\text{H}} + \text{tr} \left(\log \left(I + \frac{\eta}{1 - \eta} X \right) \right).$$

585 Now we are able to prove Theorem 3.3.

586 *Proof of Theorem 3.3.* By Lemma C.1 and the definition that $\rho_t = W_t / \text{tr}(W_t)$, we write

$$\begin{aligned} & -\log \langle A_t, \rho_t \rangle_{\text{H}} - (-\log \langle A_t, \rho \rangle_{\text{H}}) \\ &= \log \left\langle \frac{A_t}{\langle A_t, \rho_t \rangle_{\text{H}}}, \rho \right\rangle_{\text{H}} \\ &\leq \frac{1}{\eta} \left\langle \log \left((1 - \eta)I + \eta \frac{A_t}{\langle A_t, \rho_t \rangle_{\text{H}}} \right), \rho \right\rangle_{\text{H}} + \text{tr} \left(\log \left(I + \frac{\eta}{1 - \eta} \frac{A_t}{\langle A_t, \rho_t \rangle_{\text{H}}} \right) \right) \\ &= \frac{1}{\eta} \left\langle \log \left((1 - \eta)I + \eta \text{tr}(W_t) \frac{A_t}{\langle A_t, W_t \rangle_{\text{H}}} \right), \rho \right\rangle_{\text{H}} + \text{tr} \left(\log \left(I + \frac{\eta \text{tr}(W_t)}{1 - \eta} \frac{A_t}{\langle A_t, W_t \rangle_{\text{H}}} \right) \right), \end{aligned}$$

587 for all density matrix ρ such that $\langle A_t, \rho \rangle_{\text{H}} > 0$. Recall that Proposition 3.2 ensures $\text{tr}(W_t) \leq 1$.
 588 Then, by the iteration rule of Q-Soft-Bayes, we write

$$\begin{aligned} & -\log \langle A_t, \rho_t \rangle_{\text{H}} - (-\log \langle A_t, \rho \rangle_{\text{H}}) \\ &\leq \frac{1}{\eta} \left\langle \log \left((1 - \eta)I + \eta \frac{A_t}{\langle A_t, W_t \rangle_{\text{H}}} \right), \rho \right\rangle_{\text{H}} + \text{tr} \left(\log \left(I + \frac{\eta}{1 - \eta} \frac{A_t}{\langle A_t, W_t \rangle_{\text{H}}} \right) \right) \\ &= \frac{1}{\eta} \left\langle \log \left((1 - \eta)I + \eta \frac{A_t}{\langle A_t, W_t \rangle_{\text{H}}} \right), \rho \right\rangle_{\text{H}} + \text{tr} \left(\log \left((1 - \eta)I + \eta \frac{A_t}{\langle A_t, W_t \rangle_{\text{H}}} \right) - \log(1 - \eta)I \right) \\ &= \frac{1}{\eta} \langle \log(W_{t+1}) - \log(W_t), \rho \rangle_{\text{H}} + \text{tr}(\log(W_{t+1}) - \log(W_t)) - D \log(1 - \eta). \end{aligned}$$

589 By Lemma 13 in [51], we have

$$-\log(1 - \eta) \leq \bar{\eta} := \eta / (1 - \eta).$$

590 Hence, we get

$$-\log \langle A_t, \rho_t \rangle_{\text{H}} - (-\log \langle A_t, \rho \rangle_{\text{H}}) \leq \frac{1}{\eta} \langle \log(W_{t+1}) - \log(W_t), \rho \rangle_{\text{H}} + \text{tr}(\log(W_{t+1}) - \log(W_t)) + \bar{\eta}D.$$

591 Telescoping the inequality above, we write

$$\begin{aligned} & \sum_{t=1}^T (-\log \langle A_t, \rho_t \rangle_{\text{H}}) - \sum_{t=1}^T (-\log \langle A_t, \rho \rangle_{\text{H}}) \\ &\leq \frac{1}{\eta} \langle \log(W_{T+1}) - \log(W_1), \rho \rangle_{\text{H}} + \text{tr}(\log(W_{T+1}) - \log(W_1)) + \bar{\eta}TD, \end{aligned}$$

592 which indeed provides an upper bound on R_T as it holds for all possible $(A_t)_{t \in \mathbb{N}}$. Notice that
 593 $W_1 = \rho_1 = I/D$ and again, by Proposition 3.2, we have

$$\log(W_{T+1}) = \log(\rho_{T+1} \text{tr}(W_{T+1})) \leq \log(\rho_{T+1}).$$

594 Then, we write

$$\begin{aligned}
R_T &\leq \frac{1}{\eta} \langle \log(\rho_{T+1}) - \log(\rho_1), \rho \rangle_{\text{H}} + \text{tr}(\log(\rho_{T+1}) - \log(\rho_1)) + \bar{\eta}TD \\
&= \frac{1}{\eta} \langle \log(\rho_{T+1}) - \log(\rho_1), \rho \rangle_{\text{H}} + D \langle \log(\rho_{T+1}) - \log(\rho_1), \rho_1 \rangle_{\text{H}} + \bar{\eta}TD \\
&= \frac{1}{\eta} (S(\rho\|\rho_1) - S(\rho\|\rho_{T+1})) - DS(\rho_1\|\rho_{T+1}) + \bar{\eta}TD,
\end{aligned}$$

595 where S denotes the quantum relative entropy, i.e.,

$$S(\rho\|\sigma) := \langle \log(\rho) - \log(\sigma), \rho \rangle_{\text{H}}.$$

596 It is known that the quantum relative entropy is non-negative and $S(\rho\|I/D) \leq \log D$ for all $\rho \in \mathcal{D}$.

597 Therefore, we obtain

$$R_T \leq \frac{1}{\eta} \log D + \bar{\eta}TD = \frac{1}{\bar{\eta}} \log D + \bar{\eta}TD + \log D.$$

598 The regret bound is optimized when

$$\bar{\eta} = \sqrt{\frac{\log D}{TD}}.$$

599

□