Bayesian quantum state reconstruction with a learning-based tuned prior

Sangita Regmi,[1,*] Ashley N. Blackwell,[1] Amirali Khannejad,[1] Sanjaya Lohani,[1] Joseph M. Lukens,[2,3] Ryan T. Glasser,[4] Brian T. Kirby,[4,5] Thomas A. Searles[1]

[1] *Dept. of Electrical & Computer Engineering, University of Illinois Chicago, Chicago, IL 60607, USA*
[2] *Research Technology Office and Quantum Collaborative, Arizona State University, Tempe, AZ 85287, USA*
[3] *Quantum Information Science Section, Oak Ridge National Laboratory, Oak Ridge, TN 37831, USA*
[4] *Tulane University, New Orleans, LA 70118 USA*
[5] *DEVCOM Army Research Laboratory, Adelphi, MD 20783, USA*

[*] ***sangita@uic.edu***

Abstract: We demonstrate machine-learning-enhanced Bayesian quantum state tomography on near-term intermediate-scale quantum hardware. Our approach to selecting prior distributions leverages pre-trained neural networks incorporating measurement data and en-ables improved inference times over standard prior distributions.

Bayesian quantum state estimation naturally quantifies uncertainty, provides reliable estimates under any condition (e.g., any number of measurements), and minimizes mean-squared error [1]. However, these advantages come with practical computational challenges associated with calculating high-dimensional integrals, making it significantly slower than alternative methods [2]. In addition, Bayesian estimation often relies on custom likelihood functions or the direct manipulation of prior distributions to incorporate prior information. Although the first approach has shown some promise in boosting performance [3], it is conceptually unsatisfying since, in theory, all prior information ought to pass through the prior distribution. On the other hand, the current methods for directly engineering prior distributions to include prior information are relatively coarse, rely on assumptions about a system that may change over time (like the expected rank of output states), and necessitate manual tuning [4,5].

Here we experimentally implement a method, depicted in Fig. 1, that uses machine learning (ML) to define prior distributions that automatically adapt to input data and thereby address, in part, these practical and conceptual difficulties of Bayesian state reconstruction. In particular, our system uses pre-trained neural networks to automatically incorporate properties of the measured data set into the prior distribution. In addition to our system's "case-by-case" automatic tuning, we can further manually tune broad features of the prior distribution, making our approach sufficiently general to encompass several previous manual tuning techniques. Using a dataset obtained from near-term intermediate-scale (NISQ) hardware, we observe that our method reduces the net time required to perform high-fidelity Bayesian inference compared to inference with standard prior distributions.

To demonstrate the proof of concept in realistic experimental scenarios, we utilize data sets consisting of tomographic measurements performed on random quantum states implemented on *ibmq_jakarta*, one of the IBMQ Falcon processors. We first numerically generate 200 Haar-random two-qubit pure states and initialize these on *ibmq_jakarta*. Then, the states are automatically transpiled from the backend into the required quantum circuits for generation. The depths of the transpiled quantum circuits—i.e., the longest path from input to output—range between 12 and 16 gates. For each state, we perform full state tomography with a total of 36 measurement projections, corresponding to the four outcomes for all 9 two-qubit combinations of the Pauli operators $\{X,Y,Z\}_1 \otimes \{X,Y,Z\}_2$.

Our ML-based prior distribution selection method relies on first performing rapid inference using pre-trained

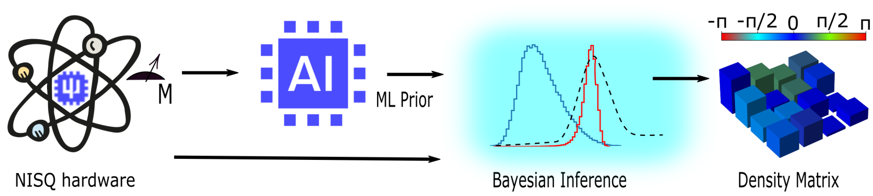

Fig. 1. Schematic of proposed machine learning prior Bayesian estimation technique.

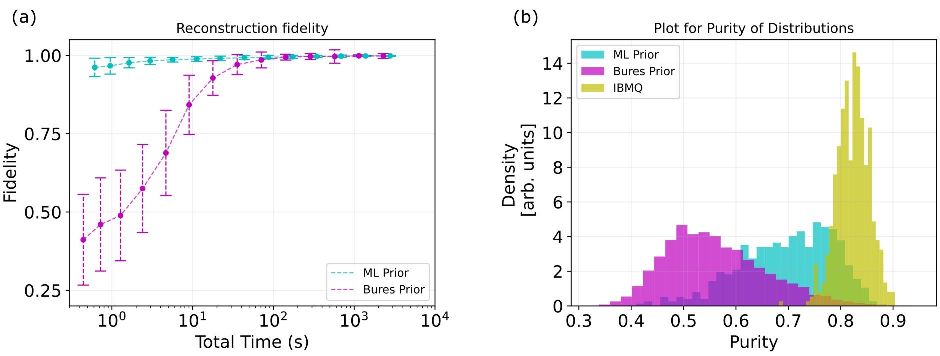

Fig. 2. Efficacy of ML-prior method for reconstructing states generated by the IBMQ machine *ibmq_jakarta*. (a) Reconstruction fidelity versus total wall-clock time. (b) Purity distributions for the ML prior, Bures prior, and IBMQ. The error bars represent one standard deviation from the mean.

neural networks. To perform this we design a convolutional neural network (CNN) with a convolutional unit of kernel size (2, 2), strides of 1, ReLU as an activation function, and 25 filters [6]. The output of the CNN is fed into a layer that performs pooling with a pool size (2, 2), followed by a second convolutional unit of the same configuration. Then, we combine two dense layers, followed by a dropout layer with a rate of 0.5, which is then attached to an output layer predicting τ-vectors (Cholesky coefficients of the density matrix). The network takes in measurements and outputs the density matrix ρ_{ML}. Note that we evaluate the fidelity of any state ρ with respect to some target density matrix ρ_0 in this paper as $\left|\mathrm{Tr}\sqrt{\sqrt{\rho_0}\rho\sqrt{\rho_0}}\right|^2$.

To find the machine-learning-defined prior (ML prior), we use the convex sum of ML-predicted density matrix ρ_{ML} and $K-1$ Haar-random pure states $|\psi_i\rangle$ as $\rho = x_1\rho_{ML} + \sum_{i=2}^{K} x_i|\psi_i\rangle\langle\psi_i|$, where, $\mathbf{x} = (x_1, \ldots, x_K)$ is a Dirichlet-distributed vector whose elements belong to the open $K-1$ simplex. The results are shown in Fig. 2 for the reconstruction of 200 randomly selected quantum states and $K = 5$. The time in these figures is wall-clock timing, meaning it includes the additional time the ML prior requires for the initial rapid inference of ρ_{ML}, a step the Bures prior does not require. The cyan and magenta curves, respectively, show the reconstruction using the ML prior and Bures prior as indicated in Fig. 2(a). The error bars represent one standard deviation from the mean. We observe that the ML prior significantly outperforms the Bures prior for short timescales, but at long computation times obtains almost the same fidelity as the Bures prior. This is expected, as our approach essentially initiates the Bayesian inference "closer" to the correct answer, but given enough time and computational resources the standard Bayesian inference approach eventually reaches the same conclusion. Furthermore, in Fig. 2(b), we show numerically generated probability density functions of purity distributions of the states from the Bures prior (magenta), the ML prior (cyan), and IBMQ experiments (yellow). We clearly observe that the ML prior, which successfully adapts based on the rapid reconstruction of ρ_{ML}, has stronger overlap in probability compared to the Bures prior distribution. Our approach provides a path toward making Bayesian quantum state estimation more practical by reducing the net computational resources required for high-fidelity inference.

Funding: U.S. Department of Energy under contract number DE-SC0012704, under contract no. DE-AC05-00OR22725. and the Early Career Research Program (ERKJ353), and the Army Research Laboratory and the Army Research Office under contract/grant numbers W911NF-19-2-0087 and W911NF-20-2-0168. The views contained in this document are those of the authors. Additionally, we acknowledge use of the IBM Quantum for this work. The views expressed are those of the authors and do not reflect the official policy or position of IBM Quantum.

References

1. C. P. Robert and G. Casella, *Monte Carlo Statistical Methods* (Springer, 1999).
2. J. M. Lukens, K. J. H. Law, and R. S. Bennink, "A Bayesian analysis of classical shadows," npj Quantum Inf. **7**, 113 (2021).
3. T. T. Mai and P. Alquier, "Pseudo-Bayesian quantum tomography with rank-adaptation," J. Stat. Plan. Inference **184**, 62–76 (2017).
4. S. Lohani, J. M. Lukens, R. T. Glasser, T. A. Searles, and B. T. Kirby, "Data-centric machine learning in quantum information science," Mach. Learn. Sci. Technol. **3**, 04LT01 (2022).
5. S. Lohani, J. M. Lukens, D. E. Jones, T. A. Searles, R. T. Glasser, and B. T. Kirby, "Improving application performance with biased distributions of quantum states," Phys. Rev. Res. **3**, 043145 (2021).
6. S. Lohani, J. M. Lukens, A. A. Davis, A. Khannejad, S. Regmi, D. E. Jones, R. T. Glasser, T. A. Searles, and B. T. Kirby, "Demonstration of machine-learning-enhanced Bayesian quantum state estimation," arXiv:2212.08032 (2022).