

Efficient Learning of Low Bond Dimension Born Machines

Maude Lizaire¹, Eric Dupuis², Guillaume Rabusseau³, Ryan Sweke⁴

¹LIS, Aix-Marseille Université, Marseille, France ²Independent, Montréal, Canada ³Mila & DIRO, Université de Montréal, Montréal, Canada ⁴AIMS, Stellenbosch University, Cape Town, South Africa

Parameterized quantum circuits (PQCs) have attracted significant interest as machine learning models, with a central open question being whether they can offer meaningful advantages over classical approaches. One natural strategy for addressing this question is to identify settings in which PQC-based models can be *dequantized*, that is, efficiently simulated or replaced by classical models without loss of performance.

Generative modelling is a particularly natural use case for PQCs: given a circuit $U(\theta)$ on n qubits, measurements yield samples from the Born distribution $p_\theta(x) = |\langle x|U(\theta)|0^{\otimes n}\rangle|^2$, allowing one to efficiently sample from complex distributions on quantum hardware. This has motivated extensive study of *Born Machines* as a route to quantum advantages in generative modelling [1]. However, despite this interest, the vast majority of existing work remains heuristic in nature, and rigorous theoretical guarantees on the learnability of Born Machines are scarce. The theoretical results that do exist are largely negative, with recent work establishing strong lower bounds on the query complexity of PAC learning Born Machines in the statistical query model [2].

In our work, we are concerned with the development of algorithms for provably dequantizing Born Machines. A natural approach to do this is to identify a *classical* model class with a similar inductive bias, and then prove that one can learn over this classical model class efficiently. To this end, we can use a key structural observation which connects Born Machines to tensor networks: Born Machines of bond dimension r can be represented as Matrix Product State (MPS) distributions of bond dimension r^2 [3]. This embedding implies that any algorithm for learning MPS distributions yields a non-proper learning algorithm for Born Machines, and motivates the study of provably efficient algorithms for MPS distributions as a path toward understanding the learnability and dequantizability of Born Machines.

In this work, we propose a spectral algorithm for learning MPS distributions that reconstructs the MPS from L -th order marginals through a TT-SVD procedure. Under a rank condition on the marginals, i.e. requiring that they share the same rank structure as the underlying MPS, the algorithm recovers the target distribution exactly. The sample and computational complexity scales linearly in the number of sites N and polynomially in the bond dimension r , avoiding the exponential dependence on N that would arise from naive approaches.

These results provide provable sufficient conditions for the efficient learnability and dequantization of Born Machines via their MPS representation. We discuss the implications of the rank condition as a characterization of the regimes in which Born Machines are efficiently learnable, and reflect on what this suggests about the potential and limitations of quantum generative models more broadly.

References

- [1] Jin-Guo Liu and Lei Wang. “Differentiable learning of quantum circuit Born machines”. In: *Physical Review A* 98.6 (Dec. 2018). ISSN: 2469-9934. DOI: 10.1103/physreva.98.062324.
- [2] Alexander Nietner et al. “On the average-case complexity of learning output distributions of quantum circuits”. In: *Quantum* 9 (Oct. 2025), p. 1883. ISSN: 2521-327X. DOI: 10.22331/q-2025-10-13-1883.
- [3] Ivan Glasser et al. *Expressive power of tensor-network factorizations for probabilistic modeling, with applications from hidden Markov models to quantum machine learning*. 2019. arXiv: 1907.03741 [cs.LG].