

Limits of quantum generative models with classical sampling hardness

Sabrina Herbst¹, Ivona Brandić¹, Adrián Pérez-Salinas²

¹Institute of Computer Engineering, TU Wien, Austria

²Institute for Theoretical Physics, ETH Zürich, Switzerland

sabrina.herbst@tuwien.ac.at

Sampling tasks have been successful in establishing quantum advantages both in theory and experiments. This has fueled the use of quantum computers for generative modeling to create samples following the probability distribution underlying a given dataset. In particular, the potential to build generative models on classically hard distributions would immediately preclude classical simulability, due to theoretical separations. Implications of using such distributions on the trainability of the models have, however, not been studied so far, which is a gap we close in this work by taking a perspective of output distributions.

In particular, we take three analytical types of probability distributions which we relate to product distributions, anticoncentrated distributions that are classically hard to sample from (e.g., IQP or Random Circuits), and peaked circuits. We study their (anti)concentration properties and implications on the concentration of commonly used loss functions as a proxy for trainability. In particular, our results show that models built on distributions, such as the ones typically found in problems showing quantum supremacy, concentrate and are therefore not trainable. In contrast, models outputting data from sparse distributions, such as peaked quantum circuits, do not suffer from exponentially concentrating losses, thus, can be trained.

We provide numerical results to showcase our analytical findings. In particular, we experiment with matrix product states, IQP circuits and peaked circuits. To extrapolate to edge cases, we additionally experiment with slightly entangled matrix product states and an analytical edge case of our anti-concentrated distribution.

Further, we discuss our results in the context of verification of quantum processes. In particular, the distributions underlying quantum supremacy experiments have been shown to require exponentially many samples for verification. Similar exponential sample complexities can be shown for statistically distinguishing distributions using the MMD². This additionally questions their use for generative models, since a bare minimum for training a quantum circuit should arguably be the possibility to efficiently verify if two probability distributions are different.

Altogether, our results show that the requirements of sampling advantage, at least in the standard forms we considered here, are incompatible with trainability. We prove that quantum advantage can still be found in generative models, though its source must be distinct from anticoncentration.

The full paper is available under [1].

References

- [1] S. Herbst, I. Brandić, and A. Pérez-Salinas. Limits of quantum generative models with classical sampling hardness, 2025. URL <https://arxiv.org/abs/2512.24801>.