Variational Quantum Generative Modeling by Sampling Expectation Values of Tunable Observables

Kevin Shen[1,2,3,*], Andrii Kurkin[1,2,3,*], Adrián Pérez-Salinas[1,4,*], Elvira Shishenina[3,*], Vedran Dunjko[1,2,*], and Hao Wang[1,2,*]

[1] $\langle aQa^L \rangle$ Applied Quantum Algorithms, Leiden University, The Netherlands
[2] LIACS, Leiden University, Niels Bohrweg 1, 2333 CA, Leiden, The Netherlands
[3] BMW Group, 80788 München, Germany
[4] Instituut-Lorentz, Leiden University, Niels Bohrweg 2, 2333 CA, The Netherlands
* This work is published in *npj Quantum Information, **11**, 178 (2025)*

Abstract. Expectation Value Samplers (EVSs) are quantum generative models that can learn high-dimensional continuous distributions by measuring the expectation values of parameterized quantum circuits. However, these models can demand impractical quantum resources for good performance. We investigate how observable choices affect EVS performance and propose an Observable-Tunable Expectation Value Sampler (OT-EVS), which achieves greater expressivity than standard EVS. By restricting the selectable observables, it is possible to use the classical shadows measurement scheme to reduce the sample complexity of our algorithm. In addition, we propose an adversarial training method adapted to the needs of OT-EVS. This training prioritizes classical updates of observables, minimizing the more costly updates of quantum circuit parameters. Numerical experiments confirm our model's expressivity and sample efficiency advantages compared to previous designs. We envision our proposal to encourage the exploration of continuous generative models running with few quantum resources.
Keywords: Generative modelling · Quantum Machine Learning · Quantum Advantage

1 Introduction

Generative modeling is the task of, given a dataset, learning to generate similar new data. Generative models such as variational autoencoders, and generative adversarial networks (GANs) have achieved remarkable success in various industrial applications, shaping diverse aspects of our daily lives. However, these models face challenges associated with high computational demands and sustainability concerns.Generative models based on quantum computers have been investigated, and provable advantages exist for specific problems. In this domain, numerous proposals have emerged, such as the Quantum Boltzmann Machine, the Quantum Circuit Born Machine, and the quantum GANs,mostly allowing for the modeling of discrete distributions.

The Expectation Value Sampler (EVS) targets continuous distribution and it is a comparatively less explored quantum generative model. In EVS, the generated data emerge as expectation values of preselected observables measured on

a state sampled from a particular distribution. Since the proposal of EVS, there have been some studies exploring its integration with classical neural networks in a hybrid framework as well as benchmarking for applications to image synthesis and molecule design.Yet, many questions regarding the theoretical properties of EVS remain unsolved.

In this work, we analyze the role of observables and propose an enhanced EVS model with tunable observables with the aim to balance quantum and classical computational resources without sacrificing performance. We provide two results in this direction. First, we show that our tunable-observable models have higher expressivity than fixed-observable models, that is the set of reachable probability distributions expands. Second, we leverage classical shadows to reduce the sample complexity of our model. To train this model, we draw inspiration from adversarial training. The key idea of the training is to train classical and quantum parameters with different update frequencies, exploiting a lower cost for classical components. We conduct numerical experiments to benchmark our adapted training method and observe its generally better performance than standard adversarial training for the same amount of quantum resources.

2 (Observable-Tunable) Expectation Value Samplers

EVSs are generative models built on parameterized quantum circuits (PQCs), designed for learning continuous distributions. We denote a expectation value sampler as a machine that transforms a latent random variable $\boldsymbol{z}$, drawn from a known and easy-to-prepare distribution, into another random variable $\boldsymbol{y}$ drawn from a distribution learnt from data. The transformation is performed via a function family $G_{\boldsymbol{\theta}}$ that applies a PQC to an initial state and measures a set of observables $\{O_m\}$ with respect to the outcoming state.

$$G_{\boldsymbol{\theta}} : \boldsymbol{z} \longmapsto \boldsymbol{y} = \left(\langle 0^{\otimes n} | U^{\dagger}_{\boldsymbol{\theta}}(\boldsymbol{z}) O_m U_{\boldsymbol{\theta}}(\boldsymbol{z}) | 0^{\otimes n} \rangle \right)_{m=1}^{M}, \tag{1}$$

An interesting property of EVSs is that the output dimension M and the number of qubits n are not intrinsically bound to each other. EVSs of only $\Theta(\log(M))$ qubits are universal approximators for M-dimensional distributions, provided $U_{\boldsymbol{\theta}}$ has infinite circuit depth.

Observable-Tunable EVS – The main contribution of this paper is the generalization of EVSs to its Observable-Tunable counterpart. The difference relies on substituting a fixed set of observables $\{O_m\}$ by the set

$$A_m = \sum_{l=1}^{L} \boldsymbol{\alpha}_{m,l} O_l, \tag{2}$$

where $\boldsymbol{\alpha}$ is a tunable matrix. For a fixed $\boldsymbol{\alpha}$, the observables $\{A_m\}$ can be used to define a standard EVS. This formulation separates the quantum (EVS) and classical (linear combination) contributions in OT-EVS, enabling precise evaluation of each component's role. In most practical settings, making the observables tunable enhances the model expressivity.

Training – The training step is central to any generative model. We choose to train train OT-EVS within the Wasserstein GAN (WGAN) framework, which

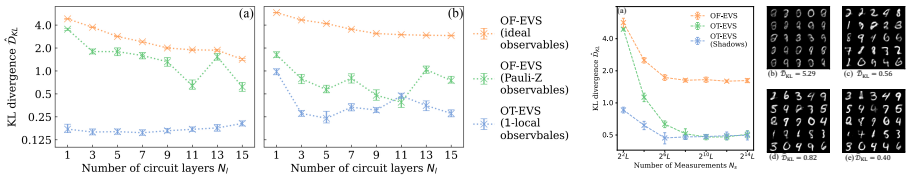


Fig. 1: **Left:** Training performance of OT-EVS and OF-EVS with varying numbers of circuit ansatz layers on synthetic datasets. (a) Sequential ansatz. (b) Brickwork ansatz. The interquartile mean and bootstraped 95% confidence intervals over 20 trials for the estimated KL divergence after $50k$ training iterations are shown. **Right:** (a) Training performance of OT-EVS with 1-local general observables and OF-EVS with 1-local Pauli-Z observables with different numbers of measurements on the MNIST dataset. The mean and standard deviation over 5 trials are shown. Samples generated by an 1-local general observable OT-EVS trained with 2^2L (b) and $2^{10}L$ (c) conventional measurements. Samples generated by an 1-local general observable OT-EVS trained with 2^2L (d) and $2^{10}L$ (e) classical shadows.

aims to minimize a distance between the generated target data distributions. The WGAN framework consists of two competing networks: a quantum generator $G_{\boldsymbol{\theta},\boldsymbol{\alpha}}$ and a classical critic $D_{\boldsymbol{w}}$ which distinguishes real and fake data samples. The two networks are trained alternately. Motivated by resource efficiency, we propose three different training schemes depending on the update frequency of $\boldsymbol{\alpha}$, $\boldsymbol{\theta}$ and $\boldsymbol{w}$. This approach reduces the consumption of quantum resources in exchange for increased classical resource usage.

Shadow-Frugal parameterization of observables – OT-EVS requires estimating the expectation values with sufficient accuracy in both the training and sampling modes. To minimize sample complexity, we propose restricting $(O_l)_{l=1}^{L}$ to maximally k-local Pauli string with $k \in \mathcal{O}(\text{polylog}(n))$, that is $L \in \Theta(n^k)$. This restriction allows for a shadow-frugal estimation of expectation values by exploiting classical shadows. To align with the training, a loss function must be estimated to an accuracy ϵ. The shadow-frugal approach allows one to estimate all observables with Pauli weight up to k with $\mathcal{O}(3^k\epsilon^{-2})$ samples, while a traditional scheme requires $\mathcal{O}(L\epsilon^{-2})$ samples.

3 Numerical experiments

We empirically evaluate OT-EVS in the shadow-frugal parameterization (see reference paper for details). All experiments use the same critic architecture for consistency. In control experiments, we observe that **a)** for a fixed quantum resource, frequent classical updating overcomes joint updates, and **b)** the shadow-frugal parametrization achieves comparable performance with fewer measurements.

We systematically compare OT-EVS and OF-EVS to assess the role of tunable observables. The training datasets comprise 8-dimensional data generated by OT-EVS models with one-layer 8-qubit circuits and 1-local general observables. For each dataset, we train models with the same circuit ansatz but various circuit depths and three observable settings: randomly-initialized tunable observables, fixed observables matching the dataset, and fixed 1-local Pauli-Z

observables. Each configuration is evaluated over 20 independent trials, with results shown in Figure 1, left. It is clear that observable-tunable models perform better than models with fixed observables. Concerning the ubiquitous problem of barren plateaus, we do not observe degrading performance as the depth increases as a consequence of step-by-step training.

We further evaluate the models on the MNIST dataset. Models are trained to learn 8-dimensional representations of the data from an autoencoder.We use a two-layer 8-qubit brickwork circuit with either tunable 1-local general observables with classical shadows or conventional measurements, or fixed 1-local Pauli-Z observables under varying measurement budgets. Model performance after 5000 training iterations (averaged over 5 independent trials) and generated samples are shown in Figure 1, right. Best performances for both models are achieved at intermediate measurement budgets, and classical shadows reduce the optimal measurement budget of the OT-EVS by at least fourfold.

4 Discussion

This work proposes a quantum generative model extending the Expectation Value Sampler using tunable observables. The model outputs random variables, given by the expectation values of an internal random state. Ideally, the outcomes follow a target probability distribution accessible through data. The tunable-observable extension enhances the expressivity of the model while reducing the quantum resources required for training, compared to fixed-observable models. These improvements are demonstrated both analytically and empirically. Our contributions are as follows. a) The expressivity enhancement is demonstrated by showing that fixed-observable models are included in tunable-observable models, not the converse. b) We can improve sample complexity by restricting our observables to those efficiently estimated using classical shadows. We also provide upper bounds on the number of measurements required to estimate the output probability distribution up to a predefined precision. c) We propose two new adversarial training procedures with improved performances while maintaining the same requirements of quantum resources, by fine-tuning the update rate of parameters in the observables and the quantum circuits.

We showcase the performance of tunable-observable models by conducting numerical experiments. First, tunable-observable models outperform fixed-observable models in practice, even in data generated by fixed-observable models. Second, the tailored training procedures and the classical-shadow sampling strategy reduce the overall training cost.

Interesting foreseeable research directions include investigating the performance of tunable-observable models on low-dimensional, real-world data sets. Technically valuable contributions are a theoretical analysis of the convergence properties and stability of the proposed training procedures and a numerical characterization of the shot noise level that benefits the model's performance.