2D and 3D Representation Learning on Gate-based Quantum Computers

Shuteng Wang Christian Theobalt Vladislav Golyanik*
MPI for Informatics, Germany

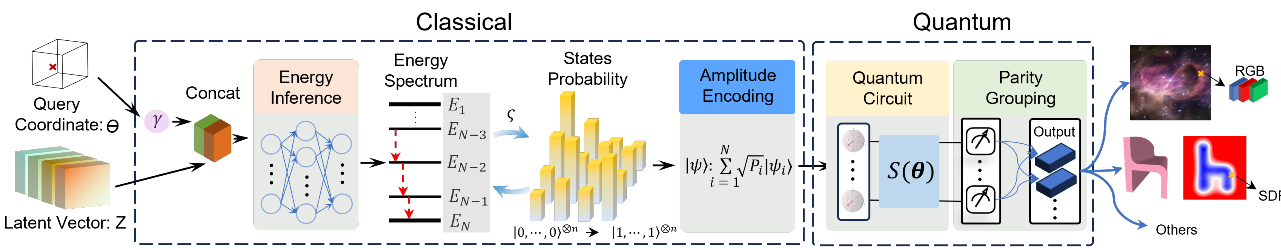

This abstract is based on a work in progress that introduces a new type of neural field for visual computing with components compatible with gate-based quantum hardware or simulators thereof. We propose a Quantum Neural Field Network (QNF-Net) expecting as input a query coordinate (of different dimensionality depending on the problem) and, optionally, a latent variable value. It outputs the corresponding learnt signal. QNF-Net includes a classical feature map providing quantum encoding of classical data and a quantum ansatz with parametrised quantum circuits.

The main figure at the top of the page provides an overview of the proposed QNF-Net. The scene coordinates $\boldsymbol{\theta}$ encoded using γ (positional encoding) concatenated with the conditioning latent code $\mathbf{z}$ are used to infer the energy spectrum $\mathbf{E}$ of a quantum system, associated with statistical uncertainty $\mathbf{P}$ modelled by, e.g., Boltzmann distribution σ. This inferred statistical property is then processed by a parametrised quantum circuitry $S(\boldsymbol{\theta})$ followed by qubit measurements. The measured values are grouped using the parity mapper to ensure consistent output dimensions.

Main Results and Implementation

We implement and test the proposed method using a high-level quantum simulator with the Pytorch interface provided in Pennylane [Bergholm et al., 2018]. We evaluate our approach for the representational accuracy of neural fields across varying data dimensions (2D and 3D). Due to the inherent computational and memory demands associated with the training of parametrised quantum circuits, we choose compact and representative data collections from CIFAR-10 [Krizhevsky and Hinton, 2009] and ShapeNet [Chang, 2015] datasets, among others.

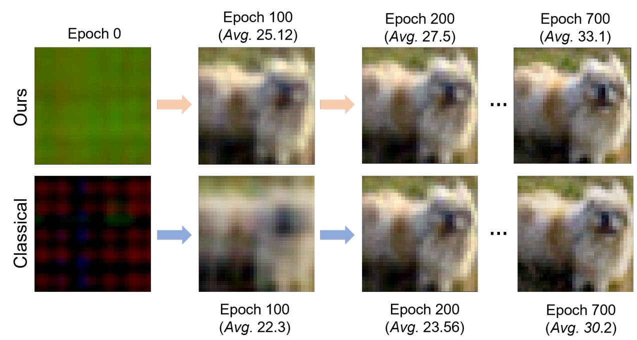

Figure 1. Visualisation of the reconstructed results for a puppy image field with increasing training epochs; mean PSNR values are provided in the brackets. Our approach (top) captures high-frequency information faster than the classical MLP (bottom).

We observe in scenarios with 2D and 3D data that the QNF-Net trained on a simulator allows us to improve both the convergence speed (see Fig. 1) and the representational accuracy (the PSNR metric emphasising high-frequency details) compared to strong classical MLP baselines. We experiment with learning 2D images of different resolutions, signed distance fields of 3D shapes as well as their collections (as learnt priors), and observe that QNF-Net substantially outperforms the classical baselines using several metrics for a comparable number of parameters in the evaluated architectures. Moreover, QNF-Net enables such applications as image or shape completion and interpolation.

In a broader sense, we practically demonstrate in a simulation the theoretically postulated advantages of quantum machine learning in the context of neural fields.

*Contact: https://4dqv.mpi-inf.mpg.de