Quantum-Assisted Variational Segmentation for Image-to-Image Wildfire Detection Using Satellite Data

Ata Akbari Asanjan[1,2], Lucas Brady[3,4], Zoe Gonzalez Izquierdo[2,4], P. Aaron Lott[2,4], Milad Memarzadeh[1,2], Nishchay Suri[2,4], David Bell[2], Eleanor Rieffel[4], Shon Grabbe[4]

[1]Data Science Group (DSG) at NASA Ames Research Center, [2]USRA Research Institute for Advanced Computer Science (RIACS), [3]KBR, [4]Quantum Artificial Intelligence Laboratory (Quail) at NASA Ames Research Center

The quantum computing community has been searching for suitable applications to demonstrate the potential of near-term quantum devices. Quantum machine learning is a potential candidate, particularly using models that cannot be efficiently simulated with classical computers [1, 2]. This work focuses on a transition phase of quantum computers where the quantum machine learning model is still simulable classically but projected not to be simulable as the size of the model grows. Ultimately quantum computers may have advantages for high-dimensional real-world problems. Due to the limited number of qubits in current noisy intermediate-scale quantum (NISQ) devices, the direct application of quantum computers in high dimensional data is not feasible. To remedy this problem, an encoder-decoder architecture can be utilized. The encoder model would transform the high-dimensional data into a compact representation, to a level that small quantum computers can be used today (or in the near future), and the decoder would take the quantum processed outputs back to the high-dimensional space.

Addressing the two challenges of quantum machine learning, this work investigates a hybrid supervised generative model with a quantum Ising Born machine embedded as the latent distribution. The model contains four main parts (Figure 1.a.): (1) a U-NET architecture responsible for learning segmentation flow, (2) a Prior network responsible for learning an encoded latent distribution of the input data, (3) a Born machine which represents the latent distribution, and (4) a Posterior network in charge of learning the joint encoded latent distribution of inputs and target data. The initial model, proposed by [3], is optimized by (1) maximizing the overlap of the prior and posterior latent distributions, and (2) minimizing the segmentation loss. The proposed model is designed to be investigated in a simulation environment applied to the real-world application of wildfire segmentation. Specifically, the model is designed to solve the patchy wildfire segmentations of Moderate Resolution Imaging Spectroradiometer (MODIS) by taking the MODIS observations and using Visible Infrared Imaging Radiometer Suite's (VIIRS) consistent wildfire product as the target. The model solves patchy wildfire segmentations and provides insight into the epistemic errors sourced from model variation. The model utilizes the Born machine as a QUBO solver to represent the latent space as a Bernoulli distribution. The proposed configuration allows the variational segmentation model to leverage the true quantum probabilistic nature and derive a more expressive latent configuration, increasing the model performance in describing wildfire segmentations. The quantum probabilistic information of the Born machine is directly incorporated in the Kullback-Leibler divergence loss in the prior and posterior distributions, forcing the Bernoulli latent distribution to maximize the overlap of input and joint input-target distributions. The proposed model is then trained and compared with a baseline only consisting of direct Bernoulli latent distribution with no Born machine representing the latent space. The models are evaluated based on the segmentation metrics, such as precision, recall, intersect of union, with uncertainty boundaries accounting for the stochastic nature of the model.

Our findings show that even in low latent-dimensional space (due to the limit in computational power of the classical quantum simulator), we are able to effectively capture the latent representation and hence the model performs better than the baseline. The findings are a projection for scaling the model into higher dimensional latent space with the Born machine surpassing the baseline performance.
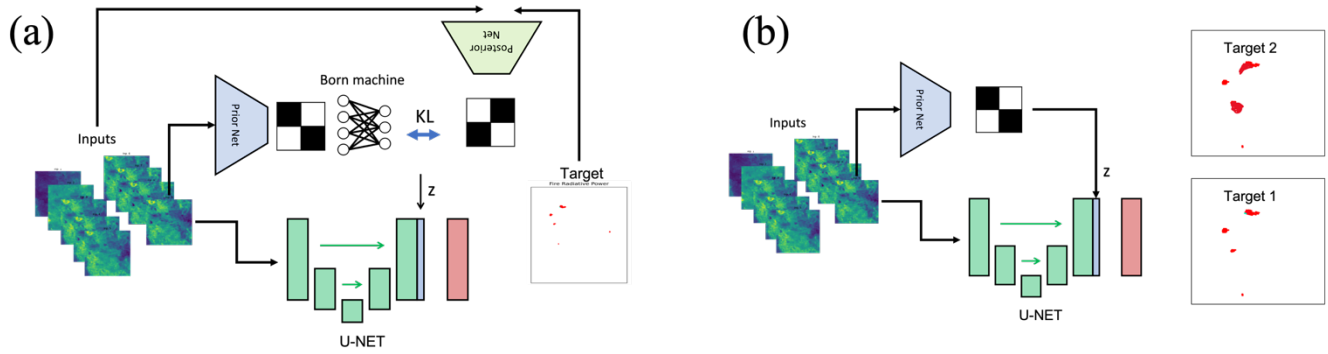
Figure 1. Sub-figure (a) demonstrates the architecture for the training phase. The model consists of a Prior and Posterior network that encode inputs and joint input-target data into compact representations, respectively. The Born machine represents the latent distribution, and the U-NET branch learns the segmentation patterns of the data. The stochasticity is introduced to the U-NET through its last layer to create meaningful but stochastic segmentations. Sub-figure (b) represents the inference phase where the model takes the stochastic behavior from the prior network and injects that into the U-NET. Each attempt of inference will generate different but similar segmentations from the same distribution of the wildfire event.

REFERENCES

[1] Coyle, B., Mills, D., Danos, V., & Kashefi, E. (2020). The Born supremacy: quantum advantage and training of an Ising Born machine. npj Quantum Information, 6(1), 1-11.

[2] Liu, J. G., & Wang, L. (2018). Differentiable learning of quantum circuit born machines. Physical Review A, 98(6), 062324.

[3] Kohl, S., Romera-Paredes, B., Meyer, C., De Fauw, J., Ledsam, J. R., Maier-Hein, K., ... & Ronneberger, O. (2018). A probabilistic u-net for segmentation of ambiguous images. Advances in neural information processing systems, 31.