Preserving Data Locality in Multidimensional Variational Quantum Classification

Mingyoung Jeng
mingyoungjeng@ku.edu
University of Kansas
Lawrence, Kansas, USA

Md. Alvir Islam Nobel
islam.alvir@ku.edu
University of Kansas
Lawrence, Kansas, USA

Vinayak Jha
vinayakjha@ku.edu
University of Kansas
Lawrence, Kansas, USA

David Levy
david.levy@ku.edu
University of Kansas
Lawrence, Kansas, USA

Dylan Kneidel
dckneidel@ku.edu
University of Kansas
Lawrence, Kansas, USA

Manu Chaudhary
manu.chaudhary@ku.edu
University of Kansas
Lawrence, Kansas, USA

SM Ishraq Ul Islam
ishraq@ku.edu
University of Kansas
Lawrence, Kansas, USA

Esam El-Araby
esam@ku.edu
University of Kansas
Lawrence, Kansas, USA

ABSTRACT

In classical machine learning, the convolution operation is leveraged in the eponymous class of convolutional neural networks (CNNs) for capturing the spatial and/or temporal locality of multidimensional input features. Preserving data locality allows CNN models to reduce the number of training parameters, and hence their training time, while achieving high classification accuracy. However, contemporary methods of quantum machine learning (QML) do not possess effective methods for exploiting data locality, due to the lack of a generalized and parameterizable implementation of quantum convolution. In this work, we propose variational quantum classification techniques that leverage a novel multidimensional quantum convolution operation with arbitrary filtering and unity stride. We provide the quantum circuits for our techniques alongside corresponding theoretical analysis. We also experimentally demonstrate the advantage of our method in comparison with existing quantum and classical techniques for image classification in staple multidimensional datasets using state-of-the-art quantum simulations.

KEYWORDS

Machine Learning, Quantum Computing, Convolutional Neural Networks

ACM Reference Format:

Mingyoung Jeng, Md. Alvir Islam Nobel, Vinayak Jha, David Levy, Dylan Kneidel, Manu Chaudhary, SM Ishraq Ul Islam, and Esam El-Araby. 2023. Preserving Data Locality in Multidimensional Variational Quantum Classification. In *Proceedings of (SC '23).* ACM, New York, NY, USA, 2 pages. https://doi.org/10.1145/nnnnnnn.nnnnnnn

1 INTRODUCTION AND BACKGROUND

Convolutional neural networks (CNNs) present an effective technique for exploiting data locality in machine learning applications such as image classification [8], which allows minimizing the number of training parameters (and thus training time) without compromising classification accuracy. However, to the best of our knowledge, contemporary quantum machine learning (QML) techniques that extend CNNs do not take full advantage of quantum computing. For example, quantum convolutional neural networks (QCNNs) [1] use variational quantum algorithms (VQAs), whose ansatz are based on CNNs, replacing neurons with qubits and quantum gates. Consequently, the resultant structure places emphasis on qubit locality rather than data locality, which is not inherently a valuable quality for capturing multidimensional features. In contrast, quanvolutional neural networks (QvNNs) [5] prefix a hybrid quantum-classical "quanvolutional" layer on a classical CNN and use a quantum circuit to perform data-local transformations. As a result, any expected quantum speedup would be nullified by the overhead of classically pre-processing data into strided windows.

In this work, we propose a novel technique of variational quantum classification. We leverage a direct, generalizable implementation of convolution and a method of pooling based on the quantum Haar transform (QHT). We experimentally compare our work to related classical and quantum algorithms using a state-of-the-art quantum simulator to show the potential of our technique.

2 PROPOSED METHODOLOGY

We propose a multidimensional quantum convolutional classifier (MQCC), which extends quantum convolution and pooling techniques to directly implement a CNN structure in a quantum circuit. Each MQCC ansatz consists of alternating trainable quantum convolution and quantum pooling layers, as well as a final trainable operation that represents the fully-connected layer.

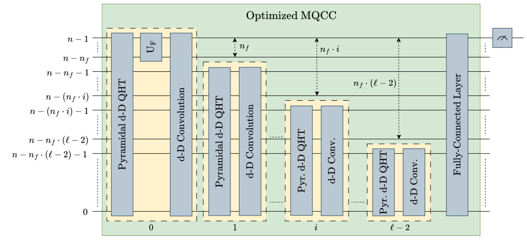

Figure 1: Optimized ℓ-layer MQCC

SC '23, NOV 12–17, 2023, Denver, CO

ACM ISBN 978-x-xxxx-xxxx-x/YY/MM…$15.00
https://doi.org/10.1145/nnnnnnn.nnnnnnn

We present a novel method of implementing convolution operations with generic multidimensional filters. The technique adds "filter" qubits and leverages quantum decrementor operations to generate shifted replicas of data. Subsequently, it embeds arbitrary multidimensional convolution filters into a unitary operation U_F, which for classification, we implement with a trainable, parameterized operation. To implement pooling, we leverage the quantum Haar transform (QHT). The Haar wavelet transform is a common method of dimension reduction in the classical domain which can be implemented easily as a quantum operation [7]. Ignoring qubit permutations, dimension reduction by a factor of 2^l on $|\psi\rangle$, where l is the number of decomposition levels and $|\psi\rangle$ is the quantum statevector representing encoded data, can be accomplished by placing parallel Hadamard gates on the l least-significant qubits of $|\psi\rangle$. In our "optimized" MQCC structure, as shown in Fig. 1, we remove the additional qubits introduced during the quantum convolution operation by applying each pooling layer before its corresponding convolution layer. However, the optimization also introduces information loss / error, which is mitigated by adding an additional trainable U_F operation. Finally, we implement the fully-connected operation as a pyramidal cascade of multiplexed rotation gates (R_y) that condenses a multi-qubit quantum circuit to a single qubit for measurement / classification.

3 EXPERIMENTAL WORK

We experimentally evaluated our proposed MQCC methods against CNNs [8], QvNNs [5], and QCNNs [1, 6] in a number of relevant metrics. Quantum simulation was performed in `Pennylane` from Xanadu [4], a framework specialized for quantum machine learning. Classical optimization was performed using `PyTorch` [3], which also provided the CNN, dataset, and loss function implementations. We configured the CNNs to match the corresponding MQCC implementation in terms of the number of convolution layers, pooling layers, and features. Experiments were performed on a high-performance computing (HPC) cluster at the University of Kansas (KU), equipped with a 48-Core Intel Xeon Gold 6342 CPU running at a base frequency of 2.8GHz, 256GB of DDR4 RAM operating at 3200 MHz, and 3× NVIDIA A100 80GB GPUs with PCIe 4.0 connectivity.

We performed 10 trials of binary classification on the MNIST [2], FashionMNIST [10], and CIFAR-10 [9] datasets at both their original resolution — (28×28), (28×28), and $(32 \times 32 \times 3)$ pixels, respectively — as well as at a down-sampled resolution of (16×16) pixels. Performance was evaluated using log-loss, training time, testing time, and testing accuracy of each technique. The quantum algorithms were also compared for their circuit depth and gate count, which are relevant to the decoherence and gate error constraints for NISQ-era quantum computers.

Our proposed MQCC (optimized) technique consistently required the fewest trainable parameters and demonstrated the lowest classification variance among all tested techniques, see Fig. 2. MQCC also demonstrated improvements in log loss, accuracy, training/testing time, and gate count across different datasets and data sizes when compared to QCNNs. In comparison to CNNs and QvNNs, MQCC also achieved faster convergence and higher average accuracy in the FashionMNIST dataset despite a higher log-loss. The high accuracy variance and low log-loss of CNNs and QvNNs suggest these models tend to overfit to the training data due to their large number of trainable parameters.

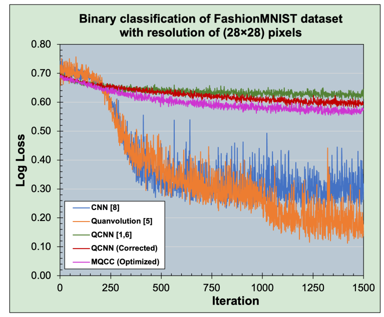

Figure 2: Log-loss of binary classification of (28×28) pixel FashionMNIST dataset

4 CONCLUSION AND FUTURE WORK

In this work, we proposed a variational quantum classification technique that preserves locality of multidimensional data. We used a direct and generalizable implementation of quantum convolution combined with quantum Haar transform (QHT)-based dimension reduction. Our experiments demonstrated improved classification performance in comparison to contemporary classical and quantum techniques. For future work, we seek to further optimize and extend our technique to multiclass classification, investigate scalability on real-world datasets, and conduct experimental trials on physical quantum hardware.

ACKNOWLEDGMENTS

This research used resources of the Oak Ridge Leadership Computing Facility, which is a DOE Office of Science User Facility supported under Contract DE-AC05-00OR22725.

REFERENCES

[1] Iris Cong, Soonwon Choi, and Mikhail D. Lukin. 2019. Quantum convolutional neural networks. *Nature Physics* 15, 12 (01 Dec 2019), 1273–1278. https://doi.org/10.1038/s41567-019-0648-8
[2] Li Deng. 2012. The mnist database of handwritten digit images for machine learning research. *IEEE Signal Processing Magazine* 29, 6 (2012), 141–142.
[3] Adam Paszke et al. 2019. PyTorch: An Imperative Style, High-Performance Deep Learning Library. In *Advances in Neural Information Processing Systems 32*. Curran Associates, Inc., 8024–8035. http://papers.neurips.cc/paper/9015-pytorch-an-imperative-style-high-performance-deep-learning-library.pdf
[4] Ville Bergholm et al. 2022. PennyLane: Automatic differentiation of hybrid quantum-classical computations. arXiv:1811.04968 [quant-ph] https://pennylane.ai/
[5] Maxwell Henderson, Samriddhi Shakya, Shashindra Pradhan, and Tristan Cook. 2020. Quanvolutional neural networks: powering image recognition with quantum circuits. *Quantum Machine Intelligence* 2, 1 (2020), 2.
[6] Tak Hur, Leeseok Kim, and Daniel K. Park. 2022. Quantum convolutional neural network for classical data classification. *Quantum Machine Intelligence* 4, 1 (10 Feb 2022), 3. https://doi.org/10.1007/s42484-021-00061-x
[7] Mingyoung Jeng et al. 2023. Improving quantum-to-classical data decoding using optimized quantum wavelet transform. *The Journal of Supercomputing* (16 Jun 2023). https://doi.org/10.1007/s11227-023-05433-7
[8] Yann LeCun, Yoshua Bengio, and Geoffrey Hinton. 2015. Deep learning. *Nature* 521, 7553 (2015), 436–444.
[9] Learning Multiple Layers of Features from Tiny Images. 2009. Alex Krizhevsky. https://www.cs.toronto.edu/~kriz/learning-features-2009-TR.pdf.
[10] Han Xiao, Kashif Rasul, and Roland Vollgraf. 2017. Fashion-MNIST: a Novel Image Dataset for Benchmarking Machine Learning Algorithms. arXiv:1708.07747 [cs.LG]