Can Quantum Computers Learn Like Classical Computers? A Co-Design Framework of Machine Learning and Quantum Circuits

Weiwen Jiang

University of Notre Dame <https://orcid.org/0000-0002-9004-487X>

Jinjun Xiong

IBM Thomas J. Watson Research Center

Yiyu Shi (✉ yshi4@nd.edu)

University of Notre Dame

Article

Keywords: quantum computers, machine learning, QF-Net, co-design framework, quantum circuit implementation, neural network

Posted Date: July 9th, 2020

DOI: <https://doi.org/10.21203/rs.3.rs-38495/v1>

License: This work is licensed under a Creative Commons Attribution 4.0 International License.
[Read Full License](#)

Version of Record: A version of this preprint was published at Nature Communications on January 25th, 2021. See the published version at <https://doi.org/10.1038/s41467-020-20729-5>.

Can Quantum Computers Learn Like Classical Computers? A Co-Design Framework for Machine Learning and Quantum Circuits

Weiwen Jiang¹, Jinjun Xiong², and Yiyu Shi¹

¹University of Notre Dame, Notre Dame, IN, 46556, USA

²IBM Thomas J. Watson Research Center, Yorktown Heights, NY, 10598, USA

ABSTRACT

Despite the pursuit of quantum supremacy in various applications, the power of quantum computers in machine learning (such as neural network models) has mostly remained unknown, primarily due to a missing link that effectively designs a neural network model suitable for quantum circuit implementation. In this article, we present the first co-design framework, namely QuantumFlow, to fix the missing link. QuantumFlow consists of a novel quantum-friendly neural network (QF-Net) design, an automatic tool (QF-Map) to generate the quantum circuit (QF-Circ) for QF-Net, and a theoretic-based execution engine (QF-FB) to efficiently support the training of QF-Net on a classical computer. We discover that, in order to make full use of the strength of quantum representation, data in QF-Net is best modeled as random variables rather than real numbers. Moreover, instead of using the classical batch normalization (which is key to achieve high accuracy for deep neural networks), a quantum-aware batch normalization method is proposed for QF-Net. Evaluation results show that QF-Net can achieve 97.01% accuracy in distinguishing digits 3 and 6 in the widely used MNIST dataset, which is 14.55% higher than the state-of-the-art quantum-aware implementation. A case study on a binary classification application is conducted. Running on IBM Quantum processor's "ibmq_essex" backend, a neural network designed by QuantumFlow can achieve 82% accuracy. To the best of our knowledge, QuantumFlow is the first framework that co-designs both the machine learning model and its quantum circuit.

Introduction

In the past decade, deep neural networks¹⁻³ have become the mainstream machine learning models, and have achieved consistent success in numerous of Artificial Intelligence applications, such as image classification⁴⁻⁷, object detection⁸⁻¹¹, and natural language processing¹²⁻¹⁴. The key factor is the significantly improved prediction accuracy by making networks deeper, known as deep learning; however, with the growing depth of neural networks, the storage and computation requirement sharply increases¹⁵, which gradually becomes the performance bottleneck in classical computers (e.g., the well-known memory wall issues¹⁶). Among all computing platforms, the quantum computer is one of the most promising ones to address such challenges^{17,18} to act as a quantum accelerator for deep neural networks¹⁹⁻²¹. Unlike classical computers with N digit bits to represent 1 N -bit number at one time, quantum computers with M qubits can represent 2^M M -bit numbers and manipulate them at the same time²². Recently, a machine learning programming framework, TensorFlow Quantum, has been proposed for quantum computer²³; however, how to exploit the power of quantum computing in deep learning is still remained unknown.

One of the most challenging obstacles to implementing deep learning algorithms on a quantum computer is the missing link between the designs of the neural networks and the corresponding quantum circuits. The existing works sepa-

ately design neural network and quantum circuits from two directions. The first direction is to map the existing neural networks designed for classical computers to quantum circuits; for instance, recent works²⁴⁻²⁷ map McCulloch-Pitts (MCP) neurons²⁸ onto quantum circuits. Such an approach can take full use of the traditional innovations in machine learning (e.g., the stochastic gradient descent in training model), but has difficulties in consistently mapping the trained model to quantum circuits. For example, it needs a large number of qubits to realize the multiplication of real numbers. To overcome this problem, some existing works²⁴⁻²⁷ assume binary representation (i.e., "-1" and "+1") of activation, which cannot well represent data as seen in modern machine learning applications. For instance, in computer vision related applications, data in images are commonly represented as real numbers. In addition, some typical operations in machine learning algorithms cannot be implemented on quantum circuits, leading to inconsistency. For example, to enable deep learning, batch normalization is a key step in a deep neural network to improve the training speed, model performance, and stability; however, directly conducting normalization on the output qubit (say normalizing the qubit with maximum probability to probability of 100%) is equivalent to reset a qubit without measurement, which is simply impossible. In consequence, batch normalization is not applied in the existing multi-layer network implementation²⁵.

Figure 1. QuantumFlow, an end-to-end co-design framework, provides a missing link between machine learning and quantum circuit designs, which consists of: (a) four sub-components QF-Net, QF-FB, QF-Circ, QF-Map that work collaboratively to design neural network and quantum implementations; (b) a data representation of data sample (e.g., from the MNIST dataset) using random variables following a two-point distribution; (c) a quantum-friendly neural network with batch normalization.

The other direction is to design neural networks dedicated to quantum computers, like the tree tensor network (TTN)^{29,30}. Such an approach has the potential to exploit quantum advantages but suffers from scalability problems. More specifically, it lacks the efficient forward/backward propagation procedure on classical computers, leading to a 2-layer neural network to take hundreds of CPU days to train a model. As such, it is simply intolerant for training a larger network, which limits the scale of networks. The effectiveness of machine learning algorithms is based on a trained model via the forward and backward propagation on large training sets. However, it is too costly to directly train one network by applying thousands of times forward and backward propagation on quantum computers; in particular, there are limited available quantum computers for public access at the current stage. An alternative way is to run a quantum simulator on a classical computer for training the models for quantum circuits, but the time complexity of quantum simulation is $O(2^m)$, where m is the number of qubits. This significantly restricts the trainable network size for quantum circuits.

To address all the above obstacles, we claim that it is demanded to take quantum circuit implementation into consideration when designing neural networks. This paper proposes the first co-design framework, namely QuantumFlow, where four sub-components (QF-Net, QF-FB, QF-Circ, and QF-Map) work collaboratively to design a neural network and implement it to a quantum computer, as shown in Figure 1(a).

QF-Net is a novel quantum-friendly neural network. In QF-Net, we discover that to take full advantage of the quantum representation, the data in a neural network, instead of being treated as real numbers, are best modeled as random variables following a two-point distribution as shown in Figure 1(b). Neural Computation (NC), one key operation in QF-Net, is designed based on such converted random variables. QF-Net also integrates a quantum-friendly batch normalization (BN) as shown in Figure 1(c). It includes additional parameters to normalize the output of a neuron, which are tuned during the training phase. To support both the infer-

ence and training of QF-Net, we further develop QF-FB, a forward/backward propagation engine for QF-Net. When QF-FB is integrated into PyTorch to conduct inference and training of QF-Net on classical computers, we denote it as QF-FB(C), whose design is based on the probability theory. QF-FB(C) is efficient for both inference and training. QF-FB can also be executed on a quantum computer or a quantum simulator. Based on Qiskit Aer simulator, we implement QF-FB(Q) for inference with or without error models.

For each operation in QF-Net (e.g., neural computation and batch normalization), a corresponding quantum circuit is designed in QF-Circ. In neural computation, an encoder is involved to encode the inputs and weights. The output will be sent to the batch normalization which involves additional control qubits to adjust the probability of a given qubit to be ranged from 0 to 1. Based on QF-Net and QF-Circ, QF-Map is an automatic tool to conduct (1) network-to-circuit mapping (from QF-Net to QF-Circ); (2) virtual-to-physic mapping (from virtual qubits in QF-Circ to physic qubits in quantum processors). Network-to-circuit mapping guarantees the consistency between QF-Net and QF-Circ with or without internal measurement; while virtual-to-physic mapping is based on Qiskit with the consideration of error rates.

As a whole, given a dataset, QuantumFlow can design and train a quantum-friendly machine learning model and automatically generate the corresponding quantum circuit. The proposed co-design framework is evaluated on the IBM Qiskit Aer simulator and IBM Quantum Processors.

Results

This section presents the evaluation results of all four sub-components in QuantumFlow. We first evaluate the effectiveness of QF-Net on the commonly used MNIST dataset³¹ for the classification task. Then, we show the consistency between QF-FB(C) on classical computers and QF-FB(Q) on the Qiskit Aer simulator. We finally conduct an end-to-end case study on a binary classification test case on IBM quantum processors to test QF-Circ and QF-Map.

Figure 2. QF-Net achieves state-of-the-art accuracy in image classifications on sub-datasets of MNIST: $\{3,6\}$, $\{3,8\}$, $\{1,3,6\}$. We employ a neural network with two hidden layers, where the first hidden layer contains 4 neurons for $\{3,6\}/\{3,8\}$ and 8 neurons for $\{1,3,6\}$; and the second hidden layer contains 2 neurons for $\{3,6\}/\{3,8\}$ and 3 neurons for $\{1,3,6\}$.

QF-Net Achieves High Accuracy on MNIST

Figure 2 reports the results of different approaches for the classification of handwritten digits on the commonly used MNIST dataset³¹. Results clearly show that with the same network structure (i.e., the same number of layers and the same number of neurons in each layer), the proposed QF-Net can achieve the highest accuracy than the existing models: (i) multi-level perceptron (MLP) with binary weights for the classical computer, denoted as MLP(C); (ii) MLP with binary inputs and weights designed for the classical computer, denoted as binMLP(C); and (iii) a state-of-the-art quantum-aware neural network with binary inputs and weights²⁵, denoted as FFNN(Q).

In the experiments, for each network, we have two implementations: one with batch normalization (w/ BN) and one without batch normalization (w/o BN). Kindly note that FFNN²⁵ does not consider batch normalization between layers. To show the benefits and generality of our newly proposed BN for improving the quantum circuits’ accuracy, we add that same functionality to FFNN for comparison. From the results, we can see that the proposed “QF-Net w/ BN” (abbr. QF-Net_BN) achieves the highest accuracy among all networks (even higher than MLP running on classical computers). Specifically, for the dataset of $\{3,6\}$, the accuracy of QF-Net_BN is 97.01%, achieving 1.84% and 14.55% higher accuracy gain against MLP(C) and FFNN(Q), respectively. Similar improvements are also achieved for QF-Net_BN on the dataset $\{3,8\}$. Because of the similarity of 3 and 8, QF-Net_BN only achieves an accuracy of 86.95%, but it is still the best accuracy among all networks. The above results validate that the proposed QF-Net has a great potential in solving machine learning problems and our co-design framework is effective to design a quantum network with high accuracy.

Furthermore, we have an interesting observation for our proposed batch normalization (BN). For all test cases, BN helps to improve the accuracy of QF-Net, and the most significant improvement is observed for dataset $\{1,3,6\}$, from less than 60% to 87.08%. Interestingly, BN also helps to improve MLP(C) accuracy significantly for dataset $\{1,3,6\}$ (from less

Table 1. Inference accuracy and efficiency comparison between QF-FB(C) and QF-FB(Q) on MNIST dataset to show the consistency of implementations of QF-Net on classical computers and quantum computers.

dataset	Qbits (Neurons)		Accuracy			Elapsed Time (CPU Time)		
	L1	L2	QF-FB(C)	QF-FB(Q)	Diff.	QF-FB(C)	QF-FB(Q)	Diff.
$\{3,6\}$	28(4)	12(2)	97.10%	95.53%	1.57%	5.13S	2,555H	1.79×10^6
$\{3,8\}$	28(4)	12(2)	86.84%	83.59%	3.25%	5.59S	2,631H	1.69×10^6
$\{1,3,6\}$	28(8)	18(3)	87.91%	81.99%	5.92%	15.89S	14,650H	3.32×10^6

than 60% to 81.99%), with a slight accuracy improvement for dataset $\{3,6\}$ and a slight accuracy drop for dataset $\{3,8\}$. This shows that the importance of batch normalization in improving model performance and the proposed BN is definitely useful for quantum neural networks.

QF-FB(C) and QF-FB(Q) are Consistent

Next, we evaluate the results of QF-FB(C) for QF-Net on classical computers, and that of QF-FB(Q) for the quantum circuit QF-Circ built upon QF-Net via QF-Map. Table 1 reports the comparison results in accuracy and elapsed time, where results under Column QF-FB(C) are the golden results. Because of the limitation of Qiskit Aer (whose backend is “ibmq_qasm_simulator”) used in QF-FB(Q) that can maximally support 32 qubits, we have to measure the results after each neuron. Specifically, in the first hidden layer, it needs 23 qubits (16 input qubits, 4 encoding qubits, and 3 auxiliary qubits) for neural computation and 4 qubits for batch normalization, and 1 output qubit; as a result, it requires 28 qubits in total. To add a new neuron, it requires the additional 8 qubits for encoding and batch normalization, which exceeds the limitation of 32 qubit for implementing two neurons on one circuit. In consequence, we repeatedly simulate each neuron in QF-Net and forward the output of neurons to the next layer. The number of qubits used for each hidden layer (“L1” and “L2”) is reported in column “Qbits”, where numbers in parenthesis indicate the number of neurons in a hidden layer.

Column “Accuracy” in Table 1 reports the accuracy comparison. We can see that there is a small difference between QF-FB(C) and QF-FB(Q). Specifically, the results obtained by QF-FB(C) have slightly higher accuracy over QF-FB(Q). This is because Qiskit Aer simulation used in QF-FB(Q) is based on the Monte Carlo method, and the output probability of different neurons may quite close for some cases, leading deviations. Another potential issue is that the trained model is based on QF-FB; however, as will be shown later, it is not practical to employ Qiskit Aer for training due to large elapsed time. The above results demonstrate QF-Net can be consistently implemented on classical and quantum computers.

Column “Elapsed Time” in Table 1 demonstrates the efficiency of QF-FB. The elapsed time is the inference time (i.e., forward propagation), used for executing all images in the test datasets, including 1968, 1983, and 3102 images for $\{3,6\}$,

Figure 3. Output probability comparison on QF-FB(C), QF-FB(Q)-ideal assuming perfect qubits, QF-FB(Q)-noise applying noise model for “ibmq_armonk” backend, and results of circuit design (“design 4”) in Figure 4(d) on “ibmq_armonk” backend on IBM quantum processor.

{3,8}, and {1,3,6}, respectively. As we can see from the table, QF-FB(Q) takes over 2,500 Hours for classifying 2 digits and 14,000 Hours for classifying 3 digits, while QF-FB(C) only takes less than 16 seconds for all datasets. The speedup is more than six orders of magnitude larger (i.e., $10^6 \times$). This verifies that QF-FB can provide an efficient forward propagation procedure to support the lengthy training of QF-Net.

In Figure 3, we further verify the accuracy of QF-FB by conducting a comparison for design 4 in Figure 4(d) on IBM quantum processor with “ibmq_armonk” backend. Kindly note that different backends are selected by QF-Map. For QF-FB(Q), we have two configurations: (1) QF-FB(Q)-ideal assuming perfect qubits; (2) QF-FB(Q)-noise with error models derived from “ibmq_armonk”. We launch either simulation or execution for respective approaches for 10 times, each of which is represented by a dot in Figure 3. We observe that the results of QF-FB(Q)-ideal are distributed around that generated by QF-FB(C) within 1% deviation; while QF-FB(Q)-noise obtains similar results of that on the IBM quantum processor. These results verify that the QF-Net on the classical computer can achieve consistent results with that of QF-Circ deployed on a quantum computer with perfect qubits.

QF-Circ and QF-Map on IBM Quantum Processor

This subsection further evaluates the efficacy of Quantum-Flow on IBM Quantum Processors. We first show the importance of quantum circuit optimization in QF-Circ to minimize the number of required qubits. Based on the optimized circuit design, we then deploy a 2-input binary classifier on IBM quantum processors.

Figure 4 demonstrates the optimization of a 2-input neuron step by step. All quantum circuits in Figures 4(a)-(d) achieve the same functionality, but with a different number of required qubits. The equivalency of all designs will be demonstrated in the **Supplementary Information**. Design 1 in Figure 4(a) is directly derived from the design methodology presented in **Methods** section. To optimize the circuit using fewer qubits, we first convert it to the circuit in Figure 4(b), denoted as design 2. Since there is only one controlled-Z gate from qubit I0 to qubit E/O, we can merge these two qubits, and obtain

Figure 4. Evaluation of the quantum circuits for a two-input neural computation, where weights are $\{-1,+1\}$: (a) design 1: original neural computation design; (b-d) three optimized designs (design 2-4), based on design 1; (e) the deviation of design 1-4 obtained from “ibmq_velencia” backend IBM quantum processor, using QF-FB(C) as golden results.

an optimized design in Figure 4(c) with 2 qubits, denoted as design 3. The circuit can be further optimized to use 1 qubit, as shown in Figure 4(d), denoted as design 4. The function f in design 4 is defined as follows:

$$f(\alpha, \beta) = 2 \cdot \arcsin(\sqrt{x + y - 2 \cdot x \cdot y}), \quad (1)$$

where $x = \sin^2 \frac{\alpha}{2}$, $y = \sin^2 \frac{\beta}{2}$, representing input probabilities.

To compare these designs, we deploy them onto IBM Quantum Processors, where “ibmq_velencia” backend is selected by QF-Map. In the experiments, we use the results from QF-FB(C) as the golden results. Figure 4(e) reports the deviations of all designs against the golden results. The results clearly show that design 4 is more robust than others, because it uses fewer qubits in the circuit. Specifically, the deviation of design 4 against golden results is always less than 5%, while reaching up to 13% for design 1. In the following experiments, design 4 is applied in QF-Circ.

Next, we are ready to introduce the case study on an end-to-end binary classification problem as shown in Figure 5. In this case study, we train the QF-Net based on QF-FB(C). Then, the tuned parameters are applied to generate QF-Circ. Finally, QF-Map optimizes the deployment of QF-Circ to IBM quantum processor, selecting the “ibmq_essex” backend.

The classification problem is illustrated in Figure 5(a), which is a binary classification problem (two classes) with two inputs: x and y . For instance, if $x = 0.2$ and $y = 0.6$, it indicates class 0. The QF-Net, QF-Circ, and QF-Map are demonstrated in Figure 5(b)-(d). First, Figure 5(b) shows that QF-Net consists of one hidden layer with one 2-input neuron and batch normalization. The output is the probability p_0 of class 0. Specifically, an input is recognized as class 0 if $p_0 \geq 0.5$; otherwise it is identified as class 1.

Figure 5. Results of a binary classification case study on IBM quantum processor of “ibmq_essex” backend: (a) binary classification with two inputs “x” and “y”; (b) QF-Net with trained parameters; (c) QF-Circ derived from the trained QF-Net; (d) the virtual-to-physic mapping obtained by QF-Map upon “ibmq_essex” quantum processor; (e) QF-FB(C) achieves 100% accuracy; (f) QF-FB(Q) achieves 98% accuracy where 2 marked error cases having probability deviation within 0.6% ; (g) results on “ibmq_essex” using the default mapping, achieving 68% accuracy; (h) results obtained by “ibmq_essex” with the mapping in (d), achieving 82% accuracy; shots number in all tests is set as 8,192.

The quantum circuit QF-Circ of the above QF-Net is shown in Figure 5(c). The circuit is composed of three parts, (1) neural computation, (2) batch_adj in batch normalization, and (3) indiv_adj in batch normalization. The neural computation is based on design 4 as shown in Figure 4(d). The parameter of Y gate in neural computation at qbit q_0 is determined by the inputs x and y . Specifically, $f(x, y) = 2 \cdot \arcsin(\sqrt{x + y - 2 \cdot x \cdot y})$, as shown in Formula 1. Then, batch normalization is implemented in two steps, where qbits q_2 and q_4 are initialized according to the trained BN parameters. During the process, q_1 holds the intermediate results after batch_adj, and q_3 holds the final results after indiv_adj. Finally, we measure the output on qbit q_3 ¹.

After building QF-Circ, the next step is to map qbits from the designed circuit to the physic qbits on the quantum processor, and this is achieved through our QF-Map. In this experiment, QF-Map selects “ibmq_essex” as backend with its physical properties shown in Figure 5(d), where error rates of each qbit and each connection are illustrated by different colors. By following the rules as defined by QF-Map (see Method section), we obtain the physically mapped QF-Circ shown

in Figure 5(d). For example, the input q_0 is mapped to the physical qbit labeled as 4.

After QuantumFlow goes through all the steps from input data to the physic quantum processor, we can perform inference on the quantum computer. In this experiments, we test 100 combinations of inputs from $\langle x, y \rangle = \langle 0.1, 0.1 \rangle$ to $\langle x, y \rangle = \langle 1.0, 1.0 \rangle$. First, we obtain the results using QF-FB(C) as golden results and QF-FB(Q) as quantum simulation assuming perfect qbits, which are reported in Figure 5(e) and (f), achieving 100% and 98% prediction accuracy. The results verify the correctness of the proposed QF-Net. Second, the results obtained on quantum processors are shown in Figure 5(h), which achieves 82% accuracy in prediction. For comparison, in Figure 5(g), we also show the results obtained by using the default mapping algorithm in IBM Qiskit, whose accuracy is only 68%. This result demonstrates the value of QF-Map in further improving the physically achievable accuracy on a physical quantum processor with errors.

Discussion

In summary, we propose an integrated QuantumFlow framework to co-design the machine learning models and quantum circuits. A novel quantum-aware QF-Net is first designed. Then an accurate and efficient inference engine, QF-FB, is

¹A Quirk-based example of inputs 0.2 and 0.6 leading to $f(x, y) = 1.6910$ can be accessed by https://wjiang.nd.edu/quirk_0_2_0_6.html, which is accessible at 06-19-2020. The output probability of 60.3% is larger than 50%, implying the inputs belong to class 0.

proposed to enable the training of QF-Net on classical computers. Based on QF-Net and the training results, the QF-Circ can automatically generate and optimize a corresponding quantum circuit. Finally, QF-Map will map QF-Circ to a quantum processor with the consideration of error rates of qbits.

Neural computation is one key component in QuantumFlow to achieve state-of-the-art accuracy. We show that the existing quantum-aware neural network²⁵ that interprets inputs as the binary form will degrade the network accuracy as shown in Figure 2. To address this problem, QF-Net models real number inputs as random variables following a two-point distribution. Details will be introduced in the next section.

Batch normalization is another key technique in improving accuracy, particularly when the network grows deeper. This can be seen from the results in Figure 2. One main reason is that the data passing a nonlinear function y^2 will lead to outputs that are significantly shrunken to a small range around 0 for real number representation and $1/m$ for a two-point distribution representation, where m is the number of inputs. Unlike the straightforward way of doing normalization on classical computers, it is non-trivial to normalize a set of qbits. Innovations are made in QuantumFlow to design a quantum-friendly batch normalization.

We have experimentally tested the QuantumFlow on a 32-qbit Qiskit Aer simulator and a 5-qbit IBM quantum processor based on superconducting technology. We show that the proposed quantum oriented machine learning model QF-Net can obtain state-of-the-art accuracy on the MNIST dataset. It can even outperform the conventional model on a similar scale for the classical computer. For the experiments on IBM quantum processors, we demonstrate that, even with the high error rates of the current quantum processor, QF-Net can be applied to classification tasks with high accuracy.

In order to accelerate the QF-FB on classical computers to support training, we make the assumptions that the perfect qbits are used. This enables us to apply theoretic formulations to accelerate the simulation process; however, it leads to some error in predicting the outputs of its corresponding deployment on a physical quantum processor with high error rates (such as the current IBM quantum processor with error rates in the range of 10^{-2}). However, we do not deem this as a drawback of our approach, rather this is an inherent problem of the current physical implementation of quantum processors. As the error rates get smaller with future quantum processors, it will help to narrow the gap between what QF-Net predicts and what quantum processor delivers. With the innovations on reducing the error rate of physic qbits, QF-Net will achieve better results.

Methods

We are going to introduce QuantumFlow in this section. Neural computation and batch normalization are two key components in a neural network, and we will present the design and implementation of these two components in QF-Net, QF-FB, QF-Circ, and QF-Map, respectively.

Figure 6. QF-Net: (a) neural computation with four operations, R : converting a real number ranging from 0 to 1 to a random variable, C : average sum of weighted inputs, A : non-linear activation function; E : converting random variable to a real number; (b) distributions of converted random variables; (c) batch normalization with two sub-components: *i*) batch adjustment, batch_adj, *ii*) individual adjustment, indiv_adj.

QF-Net

Figure 6 demonstrates two fundamental components in the proposed quantum-friendly neural network QF-Net: (1) neural computation in Figure 6(a); and batch normalization in Figure 6(c). We discuss the details of each component as follows.

An m -input neural computation component is illustrated in 6(a), where m input data $I_0, I_1, \dots, I_{m-1}$ and m corresponding weights $w_0, w_1, \dots, w_{m-1}$ are given. Input data I_i is a real number ranging from 0 to 1, while weight w_i is a $\{-1, +1\}$ binary number. Neural computation in QF-Net is composed of 4 operations: *i*) R : this operation converts a real number p_k of input I_k to a two-point distributed random variable x_k , where $P\{x_k = -1\} = p_k$ and $P\{x_k = +1\} = 1 - p_k$, as shown in 6(b). For example, we treat the input I_0 's real value of p_0 as the probability of x_0 that outcomes -1 while $q_0 = 1 - p_0$ as the probability that outcomes $+1$. *ii*) C : this operation calculates y as the average sum of weighted inputs, where the weighted input is the product of a converted input (say x_k) and its corresponding weight (i.e., w_k). Since x_k is a two-point random variable, whose values are -1 and $+1$ and the weights are binary values of -1 and $+1$, if $w_k = -1$, $w_k \cdot x_k$ will lead to the swap of probabilities $P\{x_k = -1\}$ and $P\{x_k = +1\}$ in x_k . *iii*) A : we consider the quadratic function as the non-linear activation function in this work, and A operation outputs y^2 where y is a random variable. *iv*) E : this operation converts the random variable y^2 to 0-1 real number by taking its expectation. It will be passed to batch normalization to be further used as the input to the next layer.

Figure 6(c) illustrates the proposed batch normalization (BN) component, where it takes the output of neural computation as input. The BN component is composed of two sub-components: batch adjustment ("batch_adj") and individ-

Figure 7. QF-FB: (a-b) distribution of random variable y and y^2 in neural computation component of QF-Net; (c-e) determination of parameters, t , θ , and γ , in batch normalization component of QF-Net.

ual adjustment (“indiv_adj”). Basically, batch_adj is proposed to avoid data to be continuously shrunk to a small range (as stated in **Discussion** section). This is achieved by normalizing the probability mean of a batch of outputs to 0.5 at the training phase, as shown in Figure 7(c)-(d). In the inference phase, the output $\hat{z}$ can be computed as follows:

$$\begin{aligned} \hat{z} &= (1 - z) \times \left(\sin^2 \frac{\theta}{2}\right) + z, & \text{if } t = 0 \\ \hat{z} &= z \times \left(\sin^2 \frac{\theta}{2}\right), & \text{if } t = 1 \end{aligned} \quad (2)$$

After batch_adj, the outputs of all neurons are normalized around 0.5. In order to increase the variety of different neurons’ output for better classification, indiv_adj is proposed. It contains a trainable parameter λ and a parameter γ (see Figure 7(e)). Since different neurons have different values of λ , the variation of outputs can be controlled by λ . In the inference phase, its output $\tilde{z}$ can be calculated as follows.

$$\tilde{z} = \hat{z} \times \left(\sin^2 \frac{\gamma}{2}\right) \quad (3)$$

The determination of parameters t , θ , and γ is conducted in the training phase, which will be introduced later in QF-FB.

QF-FB

QF-FB involves both forward propagation and backward propagation. In forward propagation, all weights and parameters are determined, and we can conduct neural computation and batch normalization layer by layer. The neural computation will compute $y = \frac{\sum_i \{x_i \times w_i\}}{m}$ and y^2 , where x_i is a two-point random variable, as shown in Figure 6(b). The distributions of y and y^2 are illustrated in Figure 7(a)-(b). It is straightforward to get the expectation of y^2 by using the distribution; however, for m inputs, it involves 2^m terms (e.g., $\prod q_i$ is one term), and leads to the time complexity to be $O(2^m)$. To reduce the time complexity, QF-FB takes advantage of independence of inputs to calculate the expectation as follows:

$$\begin{aligned} E([\sum_i w_i x_i]^2) &= E(\sum_i [w_i x_i]^2 + 2 \times \sum_i \sum_{j>i} [w_i x_i w_j x_j]) \\ &= m + 2 \times \sum_i \sum_{j>i} E(w_i x_i) \times E(w_j x_j) \end{aligned} \quad (4)$$

where $E(\sum_i [w_i x_i]^2) = m$, since $[w_i x_i]^2 = 1$ and there are m inputs in total. The above formula derives the following algorithm to conduct neural computation efficiently.

Algorithm 1: QF-FB: simulate neural computation

- Input:** (1) number of inputs m ; (2) m probabilities $\langle p_0, \dots, p_{m-1} \rangle$; (3) m weights $\langle w_0, \dots, w_{m-1} \rangle$.
- Output:** expectation of y^2
1. Expectation of random variable x_i : $e_i = E(x_i) = 1 - 2 \times p_i$;
 2. Expectation of $w_i \times x_i$: $E(w_i \times x_i) = w_i \times e_i$;
 3. Sum of pair product $sum_{pp} = \sum_i \sum_{j>i} \{E(w_i \times x_i) \times E(w_j \times x_j)\}$;
 4. Expectation of y^2 : $E(y^2) = \frac{m+2 \times sum_{pp}}{m^2}$;
 5. Return $E(y^2)$;

In the above algorithm, computations in line 3 dominate the time of the whole procedure, which is composed of two loops with the size of m . Hence, the time complexity of the above algorithm is $O(m^2)$. It successfully decreases the time complexity from $O(2^m)$ to $O(m^2)$. Results in Table 1 show that the proposed QF-FB is orders of magnitude faster than a quantum simulation on classical computers.

The forward propagation for batch normalization can be efficiently implemented based on the output of the neural computation. A code snippet is given as follows.

Algorithm 2: QF-FB: simulate batch normalization

- Input:** (1) $E(y^2)$ from neural computation; (2) parameters t , θ , γ determined by training procedure.
- Output:** normalized output $\tilde{z}$
1. Initialize z : $z = E(y^2)$;
 2. Calculate $\hat{z}$ according to Formula 2;
 3. Calculate $\tilde{z}$ according to Formula 3;
 4. Return $\tilde{z}$;

For the backward propagation, we need to determine weights and parameters (e.g., θ in BN). The typically used optimization method (e.g., stochastic gradient descent³²) is applied to determine weights. In the following, we will discuss the determination of BN parameters t , θ , γ .

The batch_adj sub-component involves two parameters, t and θ . During the training phase, a batch of outputs are generated for each neuron. Details are demonstrated in Figure 7(c)-(d) with 6 outputs. In terms of the mean of outputs in a batch p_{mean} , there are two possible cases: (1) $p_{mean} \leq 0.5$ and (2) $p_{mean} > 0.5$. For the first case, t is set to 0 and $\theta = 2 \times \arcsin(\sqrt{\frac{0.5 - p_{mean}}{1 - p_{mean}}})$ can be derived from Formula 2 by setting $\hat{z}$ to 0.5; similarly, for the second case, t is set to 1 and $\theta =$

Figure 8. QF-Circ: (a) neural computation, quantum circuit designs of four operations; (b-e): batch normalization, quantum circuit designs for different cases; (b) design of “batch_adj” for the case of $t = 0$; (c) design of “batch_adj” for the case of $t = 1$; (d) design of “indiv_adj”; (e) optimized design for a specific case when $t = 1$ in “batch_adj”.

$2 \times \arcsin(\sqrt{\frac{0.5}{p_{mean}}})$. Kindly note that the training procedure will be conducted in multiple iterations of batches. As with the method for batch normalization in the conventional neural network, we employ *moving average* to record parameters. Let x_i be the parameter of x (e.g., θ) at the i^{th} iteration, and x_{cur} be the value obtained in the current iteration. For x_i , it can be calculated as $x_i = m \times x_{i-1} + (1 - m) \times x_{cur}$, where m is the momentum which is set to 0.1 by default in the experiments.

In forward propagation, the sub-module *indiv_adj* is almost the same with *batch_adj* for $t = 0$; however, the determination of its parameter γ is slightly different from θ for *batch_adj*. As shown in Figure 7(e), the initial probability of $\hat{z}$ after *batch_adj* is p_z . The basic idea of *indiv_adj* is to move $\hat{z}$ by an angle, γ . It will be conducted in three steps: (1) we move start point at p_z to point A with the probability of $(p_z/n + 0.5) \times \lambda$, where n is the batch size and λ is a trainable variable; (2) we obtain γ by moving point A to $p = 0.5$; (3) we finally move solution at p_z by the angle of γ to obtain the final result. By replacing p_{mean} by $(p_z/n + 0.5) \times \lambda$ in *batch_adj* when $t = 1$, we can calculate γ . For each batch, we calculate the mean of γ , and we also employ the *moving average* to record γ .

QF-Circ

We now discuss the corresponding circuit design for components in QF-Net. The quantum circuit of the neural computation in Figure 6(a) is illustrated in Figure 8(a); while the design of different cases of batch normalization in Figure 6(b) are illustrated in Figure 8(b)-(e). A detailed demonstration of the equivalency between QF-Circ and QF-Net can be found in the **Supplementary Information**.

For the neural computation (NC) circuit in Figure 8(a), it is composed of m input (I) qubits, and $k = \log_2 m$ encoding (E) qubits, and 1 output (O) qubit. In accordance with the operations of QF-Net, the NC circuit is composed of three parts. In the first part, the NC circuit applies m Y gates with parameter $\theta = 2 \times \arcsin(\sqrt{p_k})$ (details see the **Supplementary Information**) to initialize input qubit I_k in terms of the input real value p_k , such that the state of I_k is changed from

$|0\rangle$ to $\sqrt{q_k}|0\rangle + \sqrt{p_k}|1\rangle$. The initialized qubits represent the random variables in Figure 6(b), where $P\{I_k = |1\rangle\} = p_k$, and $P\{I_k = |0\rangle\} = q_k$. Other qubits, including encoding qubits E_m and output qubit O , are initialized as $|0\rangle$.

The second part in the NC circuit completes the average sum of weighted inputs with a quadratic activation function, which can be further divided into four steps. In step 1, the NC circuit conducts dot product of inputs and weight on qubits I_k and prepares to encode qubits in superposition. For input qubits, we place a W component to perform the multiplication. Specifically, for qubit I_k , a X gate is placed if and only if $W_k = -1$; this will swap the amplitude/probability of I_k on states $|0\rangle$ and $|1\rangle$. For all encoding qubits, we apply the Hadamard (H) gate to make them enter superposition. In step 2, the NC circuit realizes the $m - k$ encoder (e.g., $m = 4$ and $k = 2$ is a $4 - 2$ encoder). During the encoding, all terms in y and y^2 's distribution (see Figure 7(a)-(b)) will be generated. The intuition is that the k qubits in encoder can have $2^k = m$ states, and we can encode m states into k qubits. Here, the input qubits control the flip of sign for the amplitude of encoding qubits; the controlled Z gates (solid dots) are placed to indicate $|1\rangle$, and in encoding qubits the controlled not Z gates (circles) are placed to indicate $|0\rangle$. For instance, the control gate from I_{m-1} is to flip the sign of state $|00\dots 0\rangle$ (all circles). In step 3, the H gates on encoding qubits collect the results on states $|I_0 I_1 \dots I_{m-1}\rangle \otimes |00\dots 0\rangle$. Finally, step 4 applies the control gate to extract the amplitudes in states $|I_0 I_1 \dots I_{m-1}\rangle \otimes |00\dots 0\rangle$ to qubit O .

After the second part of NC circuit, the expectation information has already been integrated to output qubit O , whose state is $|O\rangle = \sqrt{1 - E(y^2)}|0\rangle + \sqrt{E(y^2)}|1\rangle$. The third part in NC circuit is to measure qubit O to obtain the output real number $E(y^2)$. Kindly note that for a multi-layer QF-Net, there is no need to have a measurement at interfaces, because the converting operation R will initialize a qubit to the state exactly the same with $|O\rangle$. In addition, the following presented batch normalization circuit will also take $|O\rangle$ as input.

Now, we discuss the implementation of the batch normalization in quantum circuits. In these circuits, three qubits are

involved: (1) qbit I for input, which can be the output of qbit O in NC circuit without measurement, or initialized using a Y gate according to the measurement of qbit O in NC circuit; (2) qbit P conveys the parameter, which is obtained via training procedure, see details in QF-FB; (3) output qbits O , which can be directly used for the next layer or be measured to convert to a real number.

Figures 8(b)-(c) show the circuit design for two cases in batch_adj. Since parameters in batch_adj are determined in the inference phase, if $t = 0$, we will adopt the circuit in Figure 8(b), otherwise, we adopt that in Figure 8(c). Then, Figure 8(d) shows the circuit for indiv_adj. We can see that circuit in Figures 8(c) and (d) are the same except the initialization of parameters, θ and γ . For circuit optimization, we can merge the above two circuits into one by changing the input parameters to $g(\theta, \gamma)$, as shown in Figure 8(e). In this circuit, $z' = z \times \sin^2 \frac{g(\theta, \gamma)}{2}$, while for applying circuits in Figures 8(c) and (d), we will have $\tilde{z} = z \times \sin^2 \frac{\theta}{2} \times \sin^2 \frac{\gamma}{2}$. To guarantee the consistent function, we can derive that $g(\theta, \gamma) = 2 \times \arcsin(\sin \frac{\theta}{2} \times \sin \frac{\gamma}{2})$.

QF-Map

QF-Map is an automatic tool to map QF-Net to the quantum processor through two steps: network-to-circuit mapping and virtual-to-physic mapping. Figure 9 illustrates an example of mapping from QF-Net to QF-Circ. QF-Net in Figure 9(a) is a neural network with 2 hidden layers, designed based on neural computation (NC) and batch normalization (BN) sub-components in Figure 6 for classical computers. In QF-Circ, we notice that the results are stored in state $|1\rangle$ of the output qubits for both NC and BN; while the initialization operation R (using Y gate) is to encode the previous results to state $|1\rangle$. As a result, NC and BN circuits in QF-Circ can directly take the output qbit from the previous circuit without measurement as input. As such, Figure 9(b) shows a corresponding network to QF-Circ, where internal data type conversions in Figure 9(a) are removed. Of course, alternatively, BN can still take a given real number (e.g., measured results from the previous circuit) as input, but it would need a Y gate to initialize qbit I .

The network-to-circuit mapping from the network in Figure 9(b) to the quantum circuit QF-Circ is demonstrated in Figure 9(c). In QF-Circ, the outputs of the first hidden layer are independent, and therefore 2 NC components in first hidden layers are based on independent inputs, initialized with the same Y gates. There is no such an independent requirement for the outputs of the second layer, and therefore, 2 NC components use the same inputs. In this example, we assume the parameter t equals 0 in 2 BN components. Finally, the results can be measured at O_4 and O_5 .

After QF-Circ is generated, the final step is to map QF-Circ to quantum processors, called virtual-to-physic mapping. In this paper, we deploy QF-Circ to various IBM quantum processors. Virtual-to-physic mapping in QF-Map has two tasks: (1) select a suitable quantum processor backend, and (2) map qbits in QF-Net to physic qbits in the selected backend.

Figure 9. Network-to-circuit mapping in QF-Map: (a) QF-Net with 2 hidden layers for classical computer where batch normalization is applied in the first layer; (b) the corresponding QF-Net without internal measurement (i.e., no data type conversion); (c) QF-Circ mapped from QF-Netin (b).

For the first task, QF-Map will *i)* check the number of qbits needed; *ii)* find the backend with the smallest number of qbit to accommodate QF-Circ; *iii)* for the backends with the same number of qbits, QF-Map will select a backend for the minimum average error rate. The second task in QF-Map is to map qbits in QF-Net to physic qbits. The mapping follows two rules: (1) the qbit in QF-Net with more gates is mapped to the physic qbit with a lower error rate; and (2) qbits in QF-Net with connections are mapped to the physic qbits with the smallest distance.

References

1. LeCun, Y., Bengio, Y. & Hinton, G. Deep learning. *nature* **521**, 436–444 (2015).
2. Goodfellow, I., Bengio, Y. & Courville, A. *Deep learning* (MIT press, 2016).
3. Szegedy, C. *et al.* Going deeper with convolutions. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, 1–9 (2015).

4. Krizhevsky, A., Sutskever, I. & Hinton, G. E. Imagenet classification with deep convolutional neural networks. In *Advances in neural information processing systems*, 1097–1105 (2012).
5. He, K., Zhang, X., Ren, S. & Sun, J. Deep residual learning for image recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, 770–778 (2016).
6. Simonyan, K. & Zisserman, A. Very deep convolutional networks for large-scale image recognition. *arXiv preprint arXiv:1409.1556* (2014).
7. Szegedy, C., Vanhoucke, V., Ioffe, S., Shlens, J. & Wojna, Z. Rethinking the inception architecture for computer vision. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, 2818–2826 (2016).
8. Lin, T.-Y. *et al.* Feature pyramid networks for object detection. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, 2117–2125 (2017).
9. Ren, S., He, K., Girshick, R. & Sun, J. Faster r-cnn: Towards real-time object detection with region proposal networks. In *Advances in neural information processing systems*, 91–99 (2015).
10. He, K., Gkioxari, G., Dollár, P. & Girshick, R. Mask r-cnn. In *Proceedings of the IEEE international conference on computer vision*, 2961–2969 (2017).
11. Ronneberger, O., Fischer, P. & Brox, T. U-net: Convolutional networks for biomedical image segmentation. In *International Conference on Medical image computing and computer-assisted intervention*, 234–241 (Springer, 2015).
12. Young, T., Hazarika, D., Poria, S. & Cambria, E. Recent trends in deep learning based natural language processing. *IEEE Comput. IntelligenCe magazine* **13**, 55–75 (2018).
13. Sak, H., Senior, A. W. & Beaufays, F. Long short-term memory recurrent neural network architectures for large scale acoustic modeling. (2014).
14. Vaswani, A. *et al.* Attention is all you need. In *Advances in neural information processing systems*, 5998–6008 (2017).
15. Xu, X. *et al.* Scaling for edge inference of deep neural networks. *Nat. Electron.* **1**, 216–222 (2018).
16. Wulf, W. A. & McKee, S. A. Hitting the memory wall: implications of the obvious. *ACM SIGARCH computer architecture news* **23**, 20–24 (1995).
17. Arute, F. *et al.* Quantum supremacy using a programmable superconducting processor. *Nature* **574**, 505–510 (2019).
18. Steffen, M., DiVincenzo, D. P., Chow, J. M., Theis, T. N. & Ketchen, M. B. Quantum computing: An ibm perspective. *IBM J. Res. Dev.* **55**, 13–1 (2011).
19. Schuld, M., Sinayskiy, I. & Petruccione, F. An introduction to quantum machine learning. *Contemp. Phys.* **56**, 172–185 (2015).
20. Bertels, K. *et al.* Quantum computer architecture: Towards full-stack quantum accelerators. *arXiv preprint arXiv:1903.09575* (2019).
21. Cai, X.-D. *et al.* Entanglement-based machine learning on a quantum computer. *Phys. review letters* **114**, 110504 (2015).
22. Nielsen, M. A. & Chuang, I. Quantum computation and quantum information (2002).
23. Broughton, M. *et al.* Tensorflow quantum: A software framework for quantum machine learning. *arXiv preprint arXiv:2003.02989* (2020).
24. Tacchino, F., Macchiavello, C., Gerace, D. & Bajoni, D. An artificial neuron implemented on an actual quantum processor. *npj Quantum Inf.* **5**, 1–8 (2019).
25. Tacchino, F. *et al.* Quantum implementation of an artificial feed-forward neural network. *arXiv preprint arXiv:1912.12486* (2019).
26. Rebentrost, P., Bromley, T. R., Weedbrook, C. & Lloyd, S. Quantum hopfield neural network. *Phys. Rev. A* **98**, 042308 (2018).
27. Schuld, M., Sinayskiy, I. & Petruccione, F. The quest for a quantum neural network. *Quantum Inf. Process.* **13**, 2567–2586 (2014).
28. McCulloch, W. S. & Pitts, W. A logical calculus of the ideas immanent in nervous activity. *The bulletin mathematical biophysics* **5**, 115–133 (1943).
29. Shi, Y.-Y., Duan, L.-M. & Vidal, G. Classical simulation of quantum many-body systems with a tree tensor network. *Phys. review a* **74**, 022320 (2006).
30. Grant, E. *et al.* Hierarchical quantum classifiers. *npj Quantum Inf.* **4**, 1–8 (2018).
31. LeCun, Y., Bottou, L., Bengio, Y. & Haffner, P. Gradient-based learning applied to document recognition. *Proc. IEEE* **86**, 2278–2324 (1998).
32. Bottou, L. Large-scale machine learning with stochastic gradient descent. In *Proceedings of COMPSTAT'2010*, 177–186 (Springer, 2010).

Acknowledgements

This work is partially supported by IBM and University of Notre Dame (IBM-ND) Quantum program.

Author contributions statement

W.J. conceived the idea and performed quantum evaluations; J.X. and Y.S. supervised the work and improved the idea and experiment design. All authors contributed to manuscript writing and discussions about the results.

Figures

Figure 1

QuantumFlow, an end-to-end co-design framework, provides a missing link between machine learning and quantum circuit designs, which consists of: (a) four sub-components QF-Net, QF-FB, QF-Circ, QF-Map that work collaboratively to design neural network and quantum implementations; (b) a data representation of data sample (e.g., from the MNIST dataset) using random variables following a two-point distribution; (c) a quantum-friendly neural network with batch normalization.

Figure 2

QF-Net achieves state-of-the-art accuracy in image classifications on sub-datasets of MNIST: {3,6}, {3,8}, {1,3,6}. We employ a neural network with two hidden layers, where the first hidden layer contains 4 neurons for {3,6}/{3,8} and 8 neurons for {1,3,6}; and the second hidden layer contains 2 neurons for {3,6}/{3,8} and 3 neurons for {1,3,6}.

Figure 3

Output probability comparison on QF-FB(C), QF-FB(Q)-ideal assuming perfect qbits, QF-FB(Q)-noise applying noise model for “ibmq_armonk” backend, and results of circuit design (“design 4”) in Figure 4(d) on “ibmq_armonk” backend on IBM quantum processor.

Figure 4

Evaluation of the quantum circuits for a two-input neural computation, where weights are $\{-1,+1\}$: (a) design 1: original neural computation design; (b-d) three optimized designs (design 2-4), based on design 1; (e) the deviation of design 1-4 obtained from “ibm_velencia” backend IBM quantum processor, using QF-FB(C) as golden results.

Figure 5

Results of a binary classification case study on IBM quantum processor of "ibmq_essex" backend: (a) binary classification with two inputs "x" and "y"; (b) QF-Net with trained parameters; (c) QF-Circ derived from the trained QF-Net; (d) the virtual-to-physic mapping obtained by QF-Map upon "ibmq_essex" quantum processor; (e) QF-FB(C) achieves 100% accuracy; (f) QF-FB(Q) achieves 98% accuracy where 2 marked error cases having probability deviation within 0.6% ; (g) results on "ibmq_essex" using the default mapping, achieving 68% accuracy; (h) results obtained by "ibmq_essex" with the mapping in (d), achieving 82% accuracy; shots number in all tests is set as 8,192.

Figure 6

QF-Net: (a) neural computation with four operations, R: converting a real number ranging from 0 to 1 to a random variable, C: average sum of weighted inputs, A: non-linear activation function; E: converting random variable to a real number; (b) distributions of converted random variables; (c) batch normalization with two sub-components: i) batch adjustment, batch_adj, ii) individual adjustment, indiv_adj.

Figure 7

QF-FB: (a-b) distribution of random variable y and y^2 in neural computation component of QF-Net; (c-e) determination of parameters, t, θ , and γ , in batch normalization component of QF-Net.

Figure 8

QF-Circ: (a) neural computation, quantum circuit designs of four operations; (b-e): batch normalization, quantum circuit designs for different cases; (b) design of “batch_adj” for the case of $t = 0$; (c) design of “batch_adj” for the case of $t = 1$; (d) design of “indiv_adj”; (e) optimized design for a specific case when $t = 1$ in “batch_adj”.

Figure 9

Network-to-circuit mapping in QF-Map: (a) QF-Net with 2 hidden layers for classical computer where batch normalization is applied in the first layer; (b) the corresponding QF-Net without internal measurement (i.e., no data type conversion); (c) QF-Circ mapped from QF-Netin (b).

Supplementary Files

This is a list of supplementary files associated with this preprint. Click to download.

- [QuantumFlowAppendix.pdf](#)