






Efficient quantum circuits for machine learning activation functions including constant T -depth ReLU

Wei Zi ^{1,2}, Siyi Wang ³, Hyunji Kim ⁴, Xiaoming Sun ^{1,2,*}, Anupam Chattopadhyay ^{3,†} and Patrick Reberntrost^{5,6,‡}

¹State Key Lab of Processors, *Institute of Computing Technology, Chinese Academy of Sciences, Beijing 100190, China*

²School of Computer Science and Technology, *University of Chinese Academy of Sciences, Beijing 100049, China*

³School of Computer Science and Engineering, *Nanyang Technological University, Nanyang 639798, Singapore*

⁴Division of IT Convergence Engineering, *Hansung University, Seoul 02876, Republic of Korea*

⁵Centre for Quantum Technologies, *National University of Singapore, 117543, Singapore*

⁶School of Computing, *National University of Singapore, 117417, Singapore*



(Received 12 April 2024; accepted 1 September 2024; published 18 October 2024)

In recent years, Quantum Machine Learning (QML) has increasingly captured the interest of researchers. Among the components in this domain, activation functions hold a fundamental and indispensable role. Our research focuses on the development of activation functions quantum circuits for integration into fault-tolerant quantum computing architectures, with an emphasis on minimizing T -depth. Specifically, we present novel implementations of ReLU and leaky ReLU activation functions, achieving constant T -depths of 4 and 8, respectively. Leveraging quantum lookup tables, we extend our exploration to other activation functions such as the sigmoid. This approach enables us to customize precision and T -depth by adjusting the number of qubits, making our results more adaptable to various application scenarios. This study represents a significant advancement towards enhancing the practicality and application of quantum machine learning.

DOI: [10.1103/PhysRevResearch.6.043048](https://doi.org/10.1103/PhysRevResearch.6.043048)

I. INTRODUCTION

Currently, Machine Learning (ML) is attracting a lot of attention, driven by its widespread adoption across various industries. Recent breakthroughs in natural language processing models, exemplified by GPT4 [1], as well as advancements in image classification and generation [2–5], highlight the rapid advancements of this field. An important building block of machine learning architectures is feed-forward neural networks [6], which are collections of layers of activation functions and linear transformations. In previous research, many different activation functions have been proposed, such as Sigmoid, Softmax, and ReLU, which are successfully utilized in ML applications.

Quantum machine learning [7] has shown progress in both adapting common machine learning algorithms to run on future quantum computers and using machine learning techniques to learn about quantum systems [8]. Some examples include Quantum Neural Networks (QNN) [9], quantum support vector machines (QSVM) [10], quantum principal component analysis (QPCA) [11], variational quantum eigensolver (VQE) [12], parameterized quantum circuits (PQC)

[13], and variations of the quantum approximate optimization algorithm (QAOA) [14].

Recently, numerous research findings on QNNs based on quantum circuit implementations have been proposed [15–19]. It turns out that the implementation of activation functions, using quantum circuits, is fundamental and crucial in QNNs. Various methods have been utilized to implement activation functions, including quantum phase estimation [20,21], Taylor series expansion [22], quantum oracle synthesis [23], and polynomial networks [24]. However, to the best of the authors' knowledge, there are no dedicated circuit designs for implementing activation functions with a focus on fault tolerance so far. In this work, we address this gap by constructing a quantum circuit using Clifford+ T gates. We specifically focus on minimizing the T -depth of the circuits, considering the high cost associated with fault-tolerant implementations of the T gate and the limitation imposed by the coherence time of the quantum device [25]. In particular, for some activation functions that have a high implementation complexity, such as the ReLU [26], approximating them using polynomials would result in high polynomial degrees [27]. We have specifically designed quantum circuits to implement them, significantly reducing the T -depth.

In particular, we propose a quantum circuit to implement the ReLU function with a constant T -depth of 4, without using ancillary qubits. If the qubit connectivity is constrained to a 2D grid, the T -depth of the circuit for the ReLU function remains unchanged. Additionally, we implement Leaky ReLU functions using a quantum circuit with a constant T -depth of 8. The results of ReLU and Leaky ReLU are summarized in Table I. Here, we have optimized both the T -depth and ancilla count to an optimal level. Ours attains a constant

*Contact author: sunxiaoming@ict.ac.cn

†Contact author: anupam@ntu.edu.sg

‡Contact author: cqtfpr@nus.edu.sg

TABLE I. The T -depth of the quantum circuit implementing ReLU in both scenarios: with and without the constraint of qubit connectivity limited to a 2D grid. The T -depth for Leaky ReLU is also presented.

Activation function	ReLU	ReLU (on 2D grid)	Leaky ReLU
T -depth	4	4	8

T -depth (independent of input size) simultaneously without the employment of any ancillary qubits. This implies that the ReLU and Leaky ReLU functions can be implemented at an extremely low cost in fault-tolerant quantum computation through our construction [25].

Although ReLU and Leaky ReLU functions have efficient dedicated quantum circuit implementation. For other activation functions, like Sigmoid, Softmax, Tanh, Swish, Exponential Linear Unit (ELU) [26], and Gaussian Error Linear Unit (GELU) [28], it is not easy to find a low-cost quantum circuit implementation. Thus, we utilize a general implementation method: Quantum Look-Up Table (QLUT) [29] to implement them. In our implementation, we considered the trade-off between the number of qubits and implementation accuracy, as well as the trade-off between T -depth and ancilla count. Specifically, we represent the input and output of these activation functions using well-known floating-point numbers with $n \in \{8, 16, 32, 64, 128\}$ bits. The QLUT-based implementation method ensures that the representation accuracy of the floating-point numbers solely determines the implementation error. Moreover, QLUT enables us to reduce the T -depth of the circuit by increasing the ancilla count. This makes our results more adaptable to various application scenarios. The detailed discussion is presented in Sec. V.

We focus on the optimization of T -depth of the quantum circuit, that represents the running time of the quantum circuit in fault-tolerant quantum computation [25]. However, even though the physical implementation of quantum error correction has progressed greatly in recent years [30–32], it is still hard to use real devices to implement our circuit fault-tolerantly. Thus, we only show the construction of the quantum circuit and analyze the performance. The analysis reveals that our utilization of QLUT to implement activation functions not only makes them suitable for a wide range of application environments but also enables them to outperform other methods in certain special cases. We also implement the software version of the circuits. The Qiskit implementation of quantum circuits for activation functions discussed in this paper is opensourced on GitHub [33].

The rest of the paper is organized as follows. Section II presents preliminaries such as activation functions. Sections III and IV focus on the quantum circuit implementations of ReLU and Leaky ReLU, respectively. Section V explains how to utilize QLUT to implement other activation functions such as the sigmoid function. Finally, Sec. VI provides a conclusive summary of the entire paper.

II. PRELIMINARY

Activation functions are integral components in machine learning and deep learning, providing the crucial non-linearity

required for effective learning. In the classical world, several activation functions are commonly used including ReLU, Leaky ReLU, Sigmoid, Softmax, Tanh, Swish, Exponential Linear Unit (ELU), and Gaussian Linear Unit (GELU) [26,28]. Specifically, The formulas for ReLU, Leaky ReLU, and sigmoid are as follows:

(i) ReLU: $f(x) = \max(x, 0)$.

(ii) Leaky ReLU ($0 < \alpha < 1$):

$$f(x) = \max(x, \alpha x) = \begin{cases} x & x \geq 0 \\ \alpha x & x < 0 \end{cases}$$

(iii) Sigmoid: $f(x) = 1/(1 + e^{-x})$.

For a quantum circuit U_f that can implement function f , it is defined as

$$U_f|x\rangle|0\rangle_o = |x\rangle|f(x)\rangle.$$

Here, $|x\rangle$ represents the quantum state that encodes the input of function f , $|0\rangle_o$ represents the output qubits used to store the result $f(x)$, and $|f(x)\rangle$ represents the quantum state that encodes the output $f(x)$. The main objective of this paper is to construct U_f for various activation functions using Clifford gates and T gates. Specifically, we utilize the quantum gates from the Clifford+ T gate set $\{\text{CNOT}, H, S, T\}$ to build the circuits. The commonly used X gate can be decomposed into four Clifford gates: $X = HSSH$. Due to the higher fault-tolerant implementation cost of the T gate compared to Clifford gates, our focus lies in minimizing the T -depth, which represents the circuit's runtime [25]. In our construction, three types of ancillary qubits may be utilized. Clean ancillary qubits are qubits initialized to $|0\rangle$ and reset to $|0\rangle$ after their use. Dirty ancillary qubits do not require the initial state but need to be reset after their use. Garbage qubits are initialized to $|0\rangle$ and do not need to be reset after use.

III. QUANTUM CIRCUIT IMPLEMENTATION OF RELU FUNCTION

Definition 1. A circuit C_n is capable of implementing the n -bit ReLU function, and is defined as

$$C_n|x\rangle|0\rangle_{n-1} = |x\rangle|\max(x, 0)\rangle,$$

where $x \in \{0, 1\}^n$ represents a fixed-point number, and the circuit C_n is composed of Clifford+ T gates.

Note that the input $x \in \{0, 1\}^n$ represents a fixed-point integer, where the most significant bit (MSB) corresponds to the sign bit. Consequently, the output $\max(x, 0)$ requires only $n - 1$ bits for storage. To achieve this, we utilize $n - 1$ output qubits initialized to $|0\rangle$. Thus, a total of $2n - 1$ qubits are required to implement the ReLU function.

To understand why this number is necessary, consider $2^{n-1} + 1$ nonpositive inputs. For all of these inputs, the output is 0. Storing this output value of 0 necessitates $n - 1$ qubits. Moreover, to maintain reversibility in the presence of $2^{n-1} + 1$ distinct inputs yielding the same output 0, an additional n qubits are needed. Consequently, $2n - 1$ qubits are both necessary and sufficient for this purpose.

To implement the ReLU function, it suffices to copy the remaining $n - 1$ input qubits to the output qubits when the sign bit is 0. This statement remains valid irrespective of whether we represent the input x using the 2's complement, the 1's complement, or the true form. Subsequently, we will




FIG. 1. The quantum circuit implementing the F_8 gate.

demonstrate the implementation of the ReLU function using two architectures. The first architecture permits the application of the CNOT gate to any two qubits, while the second architecture restricts the application of the CNOT gate to qubits that are adjacent on a two-dimensional (2D) grid.

A. Implementation of ReLU Function without Connectivity Constraints

The quantum fan-out gate F_n is defined as $F_n|a, b_1, \dots, b_n\rangle = |a, b_1 \oplus a, \dots, b_n \oplus a\rangle$. Its well-known $O(\log n)$ depth implementation has been widely utilized [34], therefore we omit the proof of the Lemma 1. For illustration, the implementation of F_8 is shown in Fig. 1.

Lemma 1. The quantum fan-out gate F_n can be implemented by a quantum circuit with depth $O(\log n)$ and size $O(n)$ without ancillary qubits, using only CNOT gates.

Note that the quantum fan-out gate F_n is composed solely of CNOT gates, resulting in a T -depth of 0.

Theorem 1. There is a quantum circuit C_n capable of implementing the n -bit ReLU function with a constant T -depth of 4 without ancillary qubits. Furthermore, the circuit depth is $O(\log n)$, and the circuit size is $O(n)$.

Proof. The input $x = x_1x_2 \dots x_n$, consisting of n bits, is stored in n qubits, with the sign bit represented by x_1 . The $n - 1$ initialized output qubits are denoted as $|0\rangle_1, |0\rangle_2, \dots, |0\rangle_{n-1}$. The construction of the quantum circuit C_n proceeds in three steps:

- (i) Firstly, apply an X gate to the qubit $|x_1\rangle$.
- (ii) The second step involves applying the $n - 1$ Toffoli gates. For the i -th Toffoli gate, the first control qubit is $|x_1\rangle$, the second control qubit is $|x_{i+1}\rangle$, and the target qubit is $|0\rangle_i$.
- (iii) Lastly, apply an X gate to the qubit $|x_1\rangle$.

The quantum circuit implementing the ReLU function with input $x \in \{0, 1\}^5$ is shown in Fig. 2.

To confirm the correctness of our circuit, let us examine the following scenarios. When $x_1 = 1$, all the output qubits remain unaffected, indicating a negative input. Conversely, when $x_1 = 0$, the $n - 1$ Toffoli gates copy the values of $|x_{i+1}\rangle$ to $|0\rangle_i$ for $i = 1, 2, \dots, n - 1$. Consequently, we effectively obtain $\max(x, 0)$.

Next, we demonstrate that the $n - 1$ Toffoli gates can be implemented using a quantum circuit with a T -depth of 4.




FIG. 2. The quantum circuit implementing the ReLU function with input $x \in \{0, 1\}^5$ includes input qubits $|x_i\rangle$ and output qubits $|0\rangle_j$.

The widely recognized quantum circuit for the Toffoli gate is depicted in Fig. 3 [35]. It is important to note that the $n - 1$ Toffoli gates in C_n share a common control qubit, $|x_1\rangle$, while the other control qubits and target qubits are all distinct. We employ the shared-control Toffoli technique proposed in [36]. For visualization, the implementation of four Toffoli gates is depicted in Fig. 4.

Note that in the circuit of the Toffoli gate, the CNOT gates with $|x_1\rangle$ as the control qubit are replaced with quantum fan-out gates in the circuit of shared-control Toffoli gates, as illustrated within the dashed box in Figs. 3 and 4. Additionally, for each Toffoli gate, a T gate is applied to $|x_1\rangle$, such as the T^4 gate shown in Fig. 4.

To analyze the circuit complexity, two key observations are noted. Firstly, as stated in Lemma 1, the quantum fan-out gate does not contribute to the T -depth. Secondly, considering that the T -depth of the T^{n-1} gate is not greater than 1 due to $T^2 = S$, where the S gate is a Clifford gate. Hence, the T -depth of our circuit implementing the ReLU function is 4. According to Lemma 1, its circuit depth and size are $O(\log n)$ and $O(n)$, respectively. ■

Now, we establish the lower bound on the depth and size of the quantum circuit implementing the ReLU function. Our proof of the lower bound is inspired by Theorem 3.2 in [37].

Theorem 2. If a quantum circuit C implements the n -bit ReLU function, then the depth and size of C are lower bounded by $\Omega(\log n)$ and $\Omega(n)$, respectively.

Proof. The size lower bound is directly derived from the fact that the output is influenced by all input bits. Regarding the depth lower bound, the key observation in our proof is that the sign bit x_1 could affect each output bit. We demonstrate that if one qubit can influence the measurement result of $n - 1$ qubits, then the circuit depth has a lower bound of $\Omega(\log n)$.

We assume the existence of a quantum circuit C that implements the ReLU function. This implies that for a ReLU




FIG. 3. The quantum circuit for the Toffoli gate.

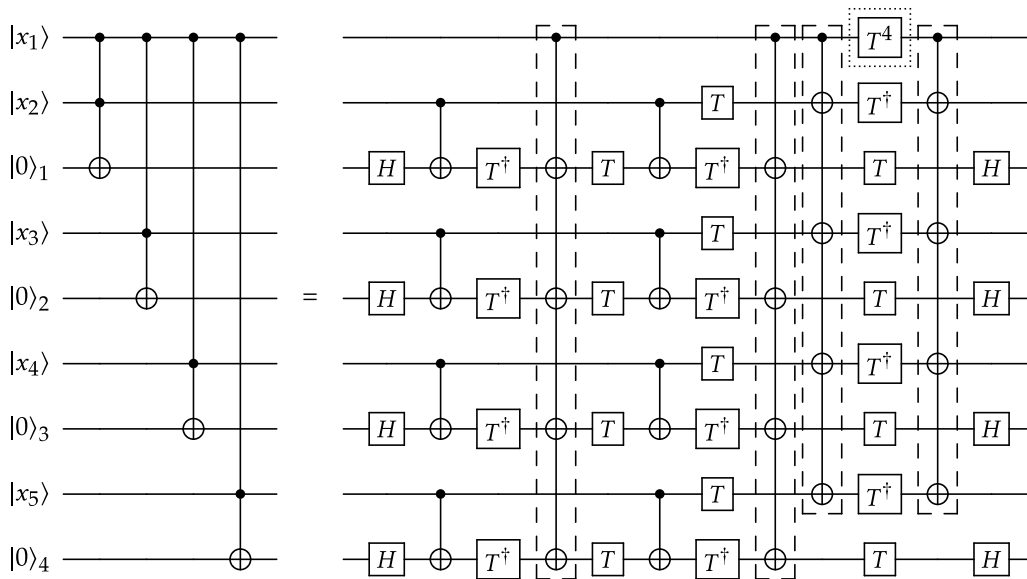


FIG. 4. The quantum circuit for the four shared-control Toffoli gates.

function R , the following equality holds:

$$\begin{aligned} &\langle x_1, \dots, x_n, 0, \dots, 0 | C^\dagger M_i C | x_1, \dots, x_n, 0, \dots, 0 \rangle \\ &= \langle x_1, \dots, x_n, 0, \dots, 0 | R^\dagger M_i R | x_1, \dots, x_n, 0, \dots, 0 \rangle. \end{aligned}$$

Here, $i \in \{1, 2, \dots, n - 1\}$ and M_i represent an observable on the i -th output qubit.

Then, we consider the number of qubits that can influence the measurement result of M_i in a d -depth quantum circuit. Let us assume that the quantum circuit C consists of d layers, and it can be represented as $C = L_1 L_2 \dots L_d$, where each L_j represents a layer of quantum gates consisting of Clifford + T gates. For each d and i , we can identify minimal subsets $L'_{1,i}, L'_{2,i}, \dots, L'_{d,i}$ such that the following equality holds:

$$\begin{aligned} &L_d^\dagger L_{d-1}^\dagger \dots L_1^\dagger M_i L_1 \dots L_{d-1} L_d \\ &= L_{d,i}^\dagger L_{d-1,i}^\dagger \dots L_{1,i}^\dagger M_i L_{1,i} \dots L_{d-1,i} L_{d,i}. \end{aligned}$$

We denote $S_{d,i}$ as the set of qubits involved in the quantum gates belonging to $L'_{d,i}$. It is important to note that the sign bit x_1 should belong to all sets $S_{d,i}$ for $i \in \{1, 2, \dots, n - 1\}$. If it does not, then the sign bit cannot influence the measurement result of M_i , which leads to a contradiction.

Let a_d denote the qubit that appears most frequently in $S_{d,1}, S_{d,2}, \dots, S_{d,n-1}$. Our aim is to prove that the number of sets $S_{d,i}$ containing a_d is no more than 2^{d-1} . We prove this statement by induction on d .

For $d = 1$, each qubit appears in at most one set. This is because there is only one layer of quantum gates and $L_{1,i}$ contains only the quantum gate applied to the i -th output qubit. The other gates are canceled, resulting in $L_{1,i}^\dagger M_i L_{1,i} = L_{1,i}^\dagger M_i L_{1,i}$.

Now, suppose our statement holds for d . We observe that there are only two ways for a qubit a to be contained in $S_{d+1,i}$.

The first way is if a is already present in $S_{d,i}$, then a will also be present in $S_{d+1,i}$ since $S_{d,i} \subset S_{d+1,i}$.

The second way is if there is a CNOT gate applied between qubits a and b belonging to L_{d+1} , and b is in $S_{d,i}$. Both a and

b appear in no more than 2^{d-1} different sets $S_{d,i}$ according to our supposition. Therefore, a appears in no more than 2^d different sets $S_{d+1,i}$.

Finally, since the sign bit needs to influence all the qubits for the output, we require $2^{d-1} \geq n - 1$. Therefore, the circuit depth d has a lower bound of $\Omega(\log n)$. ■

B. Implementing the ReLU Function on a 2D Grid

In this subsection, our focus is on implementing the ReLU function using a quantum circuit on a 2D grid. We begin by detailing the construction of the circuit, followed by establishing a lower bound for implementing ReLU on a 2D grid. Finally, we demonstrate the extension of this approach to other types of connectivity constraints.

Theorem 3. There exists a quantum circuit C_n capable of implementing the n -bit ReLU function with a constant T -depth of 4 without ancillary qubits. Furthermore, the circuit satisfies the connectivity constraint of a 2D grid, with the depth and size being $O(\sqrt{n})$ and $O(n)$, respectively.

Proof. The quantum circuit presented here is derived from the circuit outlined in Theorem 1. Our approach involves strategically positioning each qubit to facilitate the direct application of all CNOT gates, with the exception of four quantum fan-out gates. Subsequently, we devise a distinct quantum circuit consisting solely of CNOT gates to execute the quantum fan-out operation within a 2D grid. Consequently, the T -depth of our circuit remains at 4.

We assume that the number of input qubits is given by $n = k^2/2$ for some even integer k . Cases where n assumes other values can be derived from this scenario. In a 2D grid layout, we position the input qubits x_1, \dots, x_n along the odd columns, while the output qubits a_1, \dots, a_{n-1} are situated across the even columns. Specifically, for each output qubit a_i , we position x_{i+1} to its immediate left. An illustration of this configuration for $n = 32$ is depicted in Fig. 5, where the unused qubit a_0 is added for aesthetic purposes.

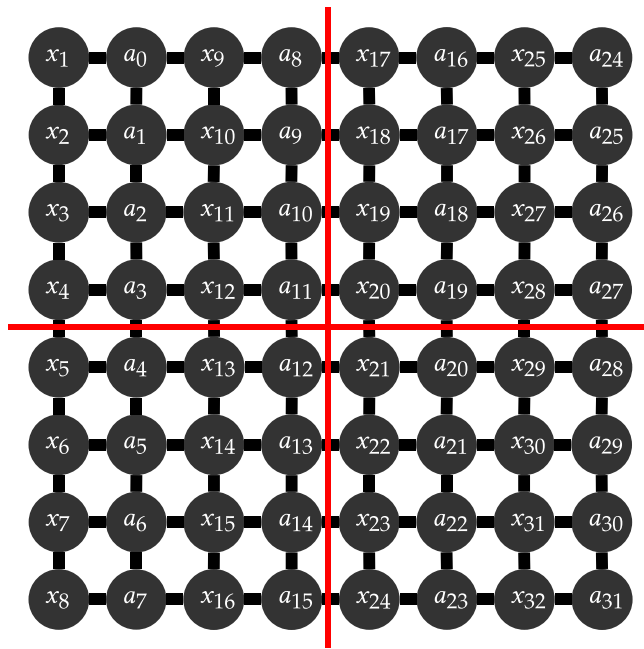


FIG. 5. An illustration of qubit arrangement on a 2D grid, with x_i representing input qubits and a_j representing output qubits. The main idea of constructing the circuit is to partition the problem into four smaller ones, as illustrated in the figure.

Note that on the 2D grid, the CNOT gate can be applied directly between two adjacent qubits. Therefore, according to Theorem 1, all the CNOT gates in the ReLU circuit can be applied directly, and our focus shifts solely to implementing the quantum fan-out gates on a 2D grid. Consider the quantum fan-out gate circuit depicted in Fig. 1. Initially, we apply a CNOT gate to add the value of $|a\rangle$ to $|b_1\rangle$. Subsequently, by utilizing CNOT gates on both sides of $|b_1\rangle$, we extend the value of $|a\rangle$ to $|b_5\rangle$. Leveraging $|b_1\rangle$ and $|b_5\rangle$, we concurrently add the value of $|a\rangle$ to $|b_3\rangle$ and $|b_7\rangle$, achieving parallelism. We adopt a similar approach on the 2D grid, dividing it evenly into four parts and transferring the value of $|x_1\rangle$ to the four smaller 2D grids. This partitions the problem into four smaller subproblems that can be processed concurrently. We then proceed with recursion. In Fig. 5, we illustrate an example where we first transfer the value of $|x_1\rangle$ to $|x_{17}\rangle$, and then concurrently transfer the value of $|x_1\rangle$ to $|x_5\rangle$ and $|x_{21}\rangle$, thus partitioning the problem into four smaller ones. To implement these long-distance CNOT gates, we utilize swap gates, which can be achieved using three CNOT gates. Denoting $D(n)$ and $S(n)$ as the depth and size of this circuit when the number of input bits is $n = k^2/2$, respectively, we establish the following recursive relation:

$$D\left(\frac{k^2}{2}\right) \leq D\left(\frac{(k/2)^2}{2}\right) + O(k), \tag{1}$$

$$S\left(\frac{k^2}{2}\right) \leq 4S\left(\frac{(k/2)^2}{2}\right) + O(k). \tag{2}$$

Then, by the master theorem [38], we have $D(n) = O(\sqrt{n})$ and $S(n) = O(n)$. We only present the quantum fan-out gate for x_i . The implementation for a_i follows a similar pattern and is omitted. If n is not equal to $k^2/2$ for some integer

k , the circuit remains similar. The crucial step is to divide the 2D grid into four smaller parts and apply recursion. This concludes the proof. ■

Theorem 4. Suppose a quantum circuit C implements the n -bit ReLU function, restricting CNOT gates in C to be applied only between nearby qubits on a 2D grid. In that case, the depth and size of C are lower bounded by $\Omega(\sqrt{n})$ and $\Omega(n)$, respectively.

Proof. The size lower bound directly follows from Theorem 2. As for the depth lower bound, similar to the proof of Theorem 2, we assume that the quantum circuit C has d layers, represented as $C = L_1 L_2 \dots L_d$, where each L_j represents a layer of quantum gates consisting of Clifford + T gates. For each d and i , we can find the minimal subset $L'_{1,i}, L'_{2,i}, \dots, L'_{d,i}$ such that the following equality holds:

$$\begin{aligned} L_d^\dagger L_{d-1}^\dagger \dots L_1^\dagger M_i L_1 \dots L_{d-1} L_d \\ = L_{d,i}^\dagger L_{d-1,i}^\dagger \dots L_{1,i}^\dagger M_i L'_{1,i} \dots L'_{d-1,i} L'_{d,i}. \end{aligned}$$

Here, $i \in \{1, 2, \dots, n-1\}$, and M_i represents an observable on the i -th output qubit. We denote $S_{d,i}$ as the set of qubits involved in the quantum gates belonging to $L'_{d,i}$.

The crucial observation is that if a qubit a is part of $S_{d+1,i}$ but not $S_{d,i}$, it indicates the presence of a CNOT gate in $L_{d+1,i}$ that acts on qubits a and b , where qubit b is in $S_{d,i}$. This implies that the minimum distance between the sign bit x_1 and all the output qubits a_1, a_2, \dots, a_{n-1} should be at most d in the 2D grid. Otherwise, $S_{d,i}$ cannot include x_1 .

Consider that within a 2D grid, the number of points within a distance of d or less is $2d(d+1)$. To ensure that this quantity is greater than or equal to $n-1$, we need $2d(d+1) \geq n-1$. Consequently, the circuit depth is lower bounded by $\Omega(\sqrt{n})$. ■

Our approach to designing a circuit for implementing the ReLU function can be extended from a 2D grid to other topological structures, such as the quantum processors from IBM. The depicted description is the following corollary.

Corollary 1. Existing quantum circuit C_n can implement the n -bit ReLU function with a constant T -depth of 4 without ancillary qubits. Furthermore, the circuit satisfies the connectivity constraint that is represented by a graph whose diameter is $O(\sqrt{n})$ with the depth and size being $O(\sqrt{n})$ and $O(n)$, respectively.

Proof. Here, the circuit satisfies the connectivity constraint represented by a graph with a diameter of $O(\sqrt{n})$. This indicates that when we employ a graph to represent which two qubits can apply the CNOT gate directly, using nodes to represent qubits and adding an edge between the corresponding two nodes if a CNOT gate can be directly applied between the two qubits. In this corollary, we stipulate that the diameter of this graph is restricted by $O(\sqrt{n})$, meaning the distance between any two nodes is limited by $O(\sqrt{n})$. Hence, a CNOT gate between any two qubits can be implemented by a circuit whose depth and size are bounded by $O(\sqrt{n})$. To construct this circuit, we can initially determine the shortest path between the two qubits, then utilize swap gates to move one qubit to the vicinity of the other to apply a CNOT gate, and finally employ swap gates to restore the state and complete the circuit construction.




FIG. 6. The qubits arrangement of 127 qubits quantum processor `ibm_brisbane` from IBM.

Note that the core of the proof of Theorem 3 lies in the formula (1). The essence of why this recursion formula holds is that a CNOT gate between any two qubits can be implemented by a circuit with a depth of $O(\sqrt{n})$. Suppose the circuit meets the connectivity constraint of a graph whose diameter is $O(\sqrt{n})$. Then, a CNOT gate between any two qubits can be implemented by a circuit with a depth of $O(\sqrt{n})$, as we have explained above. Subsequently, similar to the proof of Theorem 3, we can prove this corollary. ■

The connectivity constraint mentioned in Corollary 1 is satisfied by many quantum processors, such as the ones from IBMQ. These quantum processors have a similar connectivity structure, all consisting of cycles of 12 qubits [39]. One example is depicted in Fig. 6. Given that numerous existing quantum processors, like superconducting and neutral atom quantum processors [40,41], do not directly support the application of a CNOT gate between arbitrary two qubits, our approach for implementing the ReLU function enables it to be suitable for a significant portion of quantum processors.

IV. IMPLEMENTATION OF LEAKY RELU FUNCTION

The definition of the Leaky ReLU function is as follows:

$$f(x) = \begin{cases} x & x \geq 0 \\ \alpha x & x < 0, 0 < \alpha < 1 \end{cases}$$

Definition 2. A circuit C_n can implement the n -bit Leaky ReLU function, which means that

$$C_n|x\rangle|0\rangle_m = |x\rangle|\max(x, \alpha x)\rangle.$$

Where $x \in \{0, 1\}^n$ represents a fixed-point number, and the circuit C_n consists of Clifford+ T gates. Here, α is a predetermined parameter chosen from the set $\{0.125, 0.0625, 0.03125, 0.015625\}$. The number of output qubits m is determined by the number of bits required to store $\max(x, \alpha x)$.

Note that we choose α from the set $\{2^{-3}, 2^{-4}, 2^{-5}, 2^{-6}\}$. Therefore, $m = n + 3, n + 4, n + 5$, or $n + 6$ is sufficient to store the result $\max(x, \alpha x)$ exactly, depending on the specific choice of α . The input and output can be coded using either 2's complement or True form. The quantum circuit that implements the Leaky ReLU function will vary depending on the chosen coding method. For example, if x is coded as 1011, then $0.125x$ should be coded as 1000.011 and 1111.011 in True form and 2's complement, respectively. Note that the n and m qubits are needed to store the input and output, respectively. When the input has n bits, we will use $|0\rangle_1|0\rangle_2 \cdots |0\rangle_n$ to store the integer part and $|0\rangle_{n+1} \cdots |0\rangle_m$ to store the decimal part.

Theorem 5. There exists a quantum circuit C_n capable of implementing the n -bit Leaky ReLU function with a constant T -depth of 8 without ancillary qubits. Additionally, the circuit depth and size are $O(\log n)$ and $O(n)$, respectively.

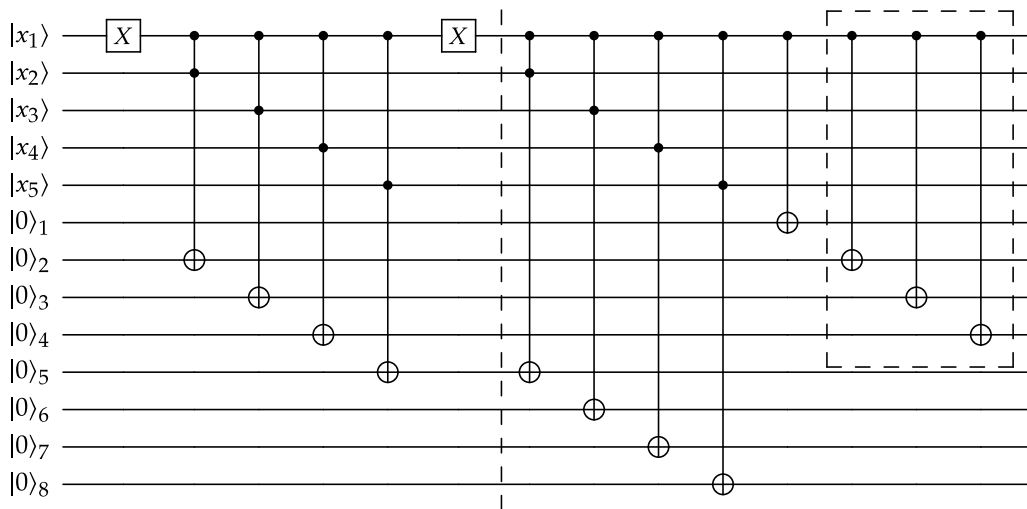


FIG. 7. The quantum circuit for the 5-bit Leaky ReLU function. Here, $|x_i\rangle$ represents input qubits, and $|0\rangle_j$ represents output qubits.

Proof. To implement the Leaky ReLU function, we output x if the sign bit is 0, and we output αx if the sign bit is 1. It is worth noting that we choose α from the set $\{0.125, 0.0625, 0.03125, 0.015625\} = \{2^{-3}, 2^{-4}, 2^{-5}, 2^{-6}\}$. Therefore, αx can be easily obtained from x by performing a bit shift. We denote the input $x = x_1 x_2 \cdots x_n$, where $x \in \{0, 1\}^n$. Furthermore, there are m output qubits denoted by $|0\rangle_1, |0\rangle_2, \dots, |0\rangle_m$. The value of m is determined by the choice of α .

The quantum circuit is constructed in two parts. The first part resembles the circuit in Theorem 1: we start by applying an X gate to $|x_1\rangle$, followed by $n - 1$ Toffoli gates. The i -th Toffoli gate has $|x_1\rangle$ as the first control qubit, $|x_{i+1}\rangle$ as the second control qubit, and $|0\rangle_{i+1}$ as the target qubit. Finally, we apply another X gate to $|x_1\rangle$. This first part of the circuit copies the state of $|x_i\rangle$ to $|0\rangle_i$ for $i = 2, 3, \dots, n$ when $x_1 = 0$, corresponding to the case where $x \geq 0$ and we should output x itself.

The second part of the circuit is constructed as follows: we apply $n - 1$ Toffoli gates, where the i -th Toffoli gate has $|x_1\rangle$ as the first control qubit, $|x_{i+1}\rangle$ as the second control qubit, and $|0\rangle_{i+k}$ as the target qubit, with $k = 1 + \log \frac{1}{\alpha}$. Then, we apply a CNOT gate with $|x_1\rangle$ as the control qubit and $|0\rangle_1$ as the target qubit. If we use the True form, the circuit is complete. However, if we use the 2's complement, we need to apply $k - 1$ additional CNOT gates, where the control qubit is $|x_1\rangle$ and the target qubits are $|0\rangle_2, |0\rangle_3, \dots, |0\rangle_k$, respectively. An example of a quantum circuit implementing the 5-bit Leaky ReLU function with $\alpha = 0.125$ is depicted in Fig. 7. In this figure, the circuit constructed in the first part is shown on the left side of the barrier. When using the True form, the gates inside the dashed box need to be omitted.

To verify the correctness of the circuit, let us examine the value of x_1 .

When $x_1 = 0$, all quantum gates in the second part can be omitted. The circuit's correctness follows from the same proof as in Theorem 1.

When $x_1 = 1$, all quantum gates in the first part can be omitted. The circuit will then copy $|x_i\rangle$ to $|0\rangle_{i+k}$ for $i = 2, 3, \dots, n$, where $k = 1 + \log \frac{1}{\alpha}$. Subsequently, we copy $|x_1\rangle$

to $|0\rangle_1$. If we use the True form, we already obtain αx . However, if we use 2's complement, we need to flip x_2, \dots, x_k . Thus, we successfully store αx on the output qubits.

To compute the T -depth of our quantum circuit, we observe that only the Toffoli gates contribute to it. By utilizing the method illustrated in Fig. 4 [36], we find that the T -depth of our circuit is 8. Similarly, we can determine that the circuit depth and size are both $O(\log n)$ and $O(n)$, respectively. ■

V. IMPLEMENTATION OF ACTIVATION FUNCTION USING QUANTUM LOOKUP TABLE

In the preceding sections, we have designed quantum circuits specifically aimed at implementing ReLU and Leaky ReLU functions. Nevertheless, for other activation functions, it is not straightforward to discover such a distinct and efficient quantum circuit implementation. Hence, we employ a general approach, namely Quantum Look-Up Table (QLUT), to implement these other activation functions. The definition of a look-up table is a table that stores the output corresponding to all possible inputs. Therefore, a look-up table is capable of representing arbitrary functions, and the representation accuracy solely depends on how the input and output are encoded. The quantum circuit implementation of a look-up table is able to transform any encoded input to the corresponding output. Rajiv Krishnakumar *et al.* [29] proposed a method to implement a Quantum Look-up Table (QLUT) designed for quantum arithmetic functions. The QLUT offers a novel approach to data access and storage within quantum algorithms, facilitating the efficient retrieval and utilization of specific values or outcomes necessary for quantum arithmetic and other computational processes. This innovation is particularly advantageous for simplifying complex computational tasks.

To utilize QLUT to implement activation functions, we initially determine the number of bits required to encode both the input and output of the function. For instance, if the input and output are encoded using n and m bits respectively, the function can be represented as $f(x) : \{0, 1\}^n \rightarrow \{0, 1\}^m$. Next, we generate a table containing all the possible inputs along

with their corresponding outputs for the function. Finally, we employ this table to construct a quantum circuit C such that:

$$C|\mathbf{x}\rangle|00\dots 0\rangle|00\dots 0\rangle = |\mathbf{x}\rangle|f(\mathbf{x})\rangle|\text{garbage}_{\mathbf{x}}\rangle,$$

where $\text{garbage}_{\mathbf{x}}$ represents the garbage qubits. We utilize QLUT to implement various activation functions, including Sigmoid, Tanh, Swish, Softmax, ELU, and GELU. The input and output of these activation functions are floating-point numbers encoded using different numbers of bits: 8 bits, 16 bits, 32 bits, 64 bits, and 128 bits. We adhere to the IEEE 754 Standard [42] for encoding floating-point numbers, except for 8 bits. For 8-bit floating-point numbers, the encoding comprises one sign bit, four exponent bits, and three mantissa bits. In this study, both the input and output consistently employ the same encoding method. Therefore, the number of bits used to encode the input and output remains consistent.

To implement a specific activation function, such as the sigmoid function, using QLUT, we first determine the number of bits to store the input and output, such as 32 bits. Subsequently, a table is generated containing 2^{32} entries, each entry representing the input and its corresponding output with 64 bits, fully describing the sigmoid function. This table is then used to construct a quantum circuit.

Note that the choice of coding method only affects the table. One is free to use fixed-point numbers to generate the table and utilize our Qiskit code to construct the quantum circuit. Next, we introduce how to construct the circuit from the table.

A. Quantum circuit implementation of quantum Lookup table

We employ the SELECTSWAP network [43] to implement the quantum Lookup Table with some minor revisions. The SELECTSWAP network comprises two parts: SELECT and SWAP. In the SELECT part, each term of the table requires a multi-controlled gate for implementation. The role of this gate is to prepare the output of each term to the output qubits only if the input state matches the input of that term. Naturally, the number of multi-controlled gates should match the number of terms in the table. If the input necessitates n bits for storage, there will be 2^n multi-controlled gates. This results in a substantial circuit depth since we need to implement 2^n n -controlled gates sequentially. By leveraging the SWAP operator, we can mitigate the circuit depth while increasing the number of qubits. Let l denote the number of swap qubits, where $0 < l < n$. For each multi-controlled gate, the 2^l outputs are prepared in $n \cdot 2^l$ qubits. Consequently, the number of multi-controlled gates can be reduced to 2^{n-l} . After applying these gates, swap gates controlled by the swap qubits ensure that the correct output is directed to the output qubits. The SELECT part has already implemented the quantum look-up table without the requirement of ancillary qubits. The SWAP part is added to it solely to reduce the depth of the quantum circuit. However, when the SWAP part is employed to decrease the circuit depth, the required number of ancillary qubits increases. Thus, users of our method can freely select the most appropriate parameter (the number of swap qubits) to make a trade-off between circuit depth and ancilla count. A further detailed explanation follows.




FIG. 8. The quantum circuit of SWAP operator for $\mathbf{x} \in \{0, 1\}^3$.

SELECT: The SELECT operator applies different sets of unitaries (here, the X gate) based on the state of the control qubits. SELECT operator iterates through all possible addresses. Let us assume the i -th term of the table contains input $\mathbf{a}_i = a_{i,1}a_{i,2}\dots a_{i,n} \in \{0, 1\}^n$, and the corresponding output $\mathbf{b}_i = b_{i,1}b_{i,2}\dots b_{i,n} \in \{0, 1\}^n$. The SELECT operator is described as follows:

$$\text{SELECT} = \sum_i |\mathbf{a}_i\rangle\langle\mathbf{a}_i| \otimes X^{\mathbf{b}_i},$$

where $X^{\mathbf{b}_i} = \otimes_j X^{b_{i,j}}$. When implementing the i -th term of the table, two key points are considered in constructing the circuit: the control condition involves the input qubits in state $|\mathbf{a}_i\rangle$, and the X gate is applied to the output qubits based on \mathbf{b}_i and the control condition. To achieve this, we can utilize the quantum circuit of the quantum fan-out gate as depicted in Fig. 1, with the top CNOT gate replaced by an n -Toffoli gate and X gates added on both sides of the n -Toffoli gate according to \mathbf{a}_i . Note that according to Fig. 3, the T -depth of the Toffoli gate is 4. Additionally, based on Lemma 7.2 of [44], the T -depth of a $(k > 3)$ -Toffoli gate is $16k - 32$.

SWAP: Assuming the number of swap qubits is l , where $0 < l < n$, there will be 2^l output registers. The function of the SWAP operator is to exchange the order of output registers so that the first register contains the quantum state of the register indexed by the state of the swap qubits. The SWAP operator is defined as follows:

$$\text{SWAP}[|\mathbf{x}\rangle \otimes_i |\phi_i\rangle_i] = |\mathbf{x}\rangle|\phi_{\mathbf{x}}\rangle_0 \otimes \dots$$

The SWAP operator comprises controlled-SWAP gates [43]. The crucial aspect is to handle the swap qubits one by one. An example of a quantum circuit with three swap qubits is depicted in Fig. 8. It is worth noting that for many controlled-SWAP gates, the T -depth of the quantum circuit implementing them can be reduced to just 4, given the condition that they share the same control qubit and have different target qubits for each SWAP gate. This reduction is achieved using the shared-control Toffoli technique [36] mentioned in Theorem 1. An example of four controlled-SWAP gates is illustrated in Fig. 9. Therefore, the T -depth of the SWAP operator is $4l$.




FIG. 9. The parallel implementation of 4 controlled swap gates.

SELECTSWAP: The SELECTSWAP circuit is a hybrid of the SELECT and SWAP circuits. Assuming the number of input qubits and output qubits is n , and the number of swap qubits is l , the SELECT operator consists of 2^{n-l} steps. In each step, 2^l n -qubit registers are prepared to store all possible outputs based on all possible states of the l swap qubits. In the SWAP operator, the register indexed by the quantum state of the swap qubits is placed in the first register. Consequently, the first register stores the desired output. The T -depth of circuit is $(2^{n-l})(16(n-l) - 32) + 4l$ according to the discussion above. Note that the garbage qubits used in the SWAP operator could be replaced by dirty ancillary qubits [43].

B. Analysis and Comparison of QLUT

Note that the T -depth of the quantum circuit for QLUT is independent of the function it needs to implement; it is solely determined by the number of input qubits and swap qubits.

Thus, the T -depth of the Sigmoid, Tanh, Swish, Softmax, ELU, and GELU functions is presented in Table II. Here, we consider both the trade-off between qubit number and implementation accuracy, and between space and time. Since errors arise solely from coding accuracy, having more input qubits leads to a lower error rate. Increasing the number of swap qubits can significantly reduce the T -depth while also increasing the number of ancillary qubits. Users can choose the most suitable parameters based on their specific situation. For instance, a quantum circuit with a T -depth of 148 and 256 ancillary qubits could be utilized to implement activation functions such as Sigmoid and Tanh, where the input is an 8-bit floating-point number. The discussion regarding the opportunity cost of ancillary qubits explains that in this case, it is worthwhile to increase the number of ancillary qubits to reduce the T -depth [45].

We illustrate the error rate of our activation function implementation and compare it with other methods. For the sake of simplicity, we define the error rate ϵ as the maximum

TABLE II. The T -depth of the quantum circuit implementing Sigmoid, Softmax, Tanh, Swish, ELU, and GELU using QLUT with varying ancilla counts (#ancilla) is shown. The number of input qubits ranges from {8, 16, 32, 64, 128}.

Number of input qubits = 8							
#ancilla	16	32	64	128	256	512	1024
T -depth	10244	4104	1548	528	148	40	28
Number of input qubits = 16							
#ancilla	64	256	1024	4096	16384	65536	2.62×10^5
T -depth	3.15×10^6	6.55×10^5	1.31×10^5	24608	4136	560	72
Number of input qubits = 32							
#ancilla	512	8192	1.31×10^5	2.1×10^6	3.36×10^7	5.37×10^8	8.59×10^9
T -depth	1.12×10^{11}	5.91×10^9	3.2×10^8	1.47×10^7	6.55×10^5	24672	624
Number of input qubits = 64							
#ancilla	16384	4.19×10^6	1.07×10^9	2.75×10^{11}	7.04×10^{13}	1.81×10^{16}	4.61×10^{18}
T -depth	6.23×10^{19}	2.07×10^{17}	6.69×10^{14}	2.06×10^{12}	5.91×10^9	1.47×10^7	24800
Number of input qubits = 128							
#ancilla	8.39×10^6	5.5×10^{11}	3.6×10^{16}	2.36×10^{21}	1.55×10^{26}	1.01×10^{31}	6.65×10^{35}
T -depth	9.14×10^{36}	1.19×10^{32}	1.51×10^{27}	1.83×10^{22}	2.07×10^{17}	2.06×10^{12}	1.47×10^7

TABLE III. This table presents the error of activation functions implemented by utilizing QLUT for different quantities of input qubits. Additionally, the degree of polynomial necessary to approximate the sigmoid function with the same error is also displayed in the table.

input qubit number	ϵ	degree of polynomial
$n = 8$	$1/2$	137
$n = 16$	$1/2^7$	297
$n = 32$	$1/2^{21}$	678
$n = 64$	$1/2^{50}$	1469
$n = 128$	$1/2^{112}$	3158

implementation error within the range of $|x| < 15/4$, as follows:

$$|\hat{f}(x) - f(x)| < \epsilon, |x| < 15/4.$$

Here, $\hat{f}(x)$ denotes the implemented function utilized to approximate $f(x)$. The error rate of using 8, 16, 32, 64, and 128-bit floating-point numbers to represent $|x| < 15/4$ is computed and shown in Table III. It should be noted that since we utilize QLUT to implement activation functions, the error rate solely depends on the number of bits employed to encode input and output. Consequently, the implementation errors for different activation functions like sigmoid and tanh are identical and presented in Table III. For the sake of comparison, we analyze the cost of implementing the sigmoid function through the polynomial approximation method. We calculated the degree of polynomial required to approximate the sigmoid function with the same error rate (by using Lemma 2 in [24]) and depicted it in Table III. A higher degree of the polynomial implies a greater quantum cost for implementing the polynomial by using methods such as Quantum Singular Value Transformation (QSVT) [46]. The table clearly shows that even with a low error rate, the required degree remains relatively high, indicating that our method holds an advantage when the function is on a small scale.

It is worth noting that if we use polynomials to approximate the ReLU function, the cost is exponentially higher than when approximating the sigmoid function, as shown by Lemma 28 in [27]. At the same time, through our construction described in Sec. III, the ReLU function can be implemented by a constant T -depth quantum circuit without ancillary qubits. This demonstrates the great advantage of our methods.

VI. CONCLUSION

This paper focuses on implementing activation functions with lower T -depth quantum circuits, yielding significant breakthroughs. Specifically, we achieve constant T -depth circuits for both ReLU and Leaky ReLU functions, regardless of the number of input qubits, n . Even when restricting the connectivity topology between qubits to a 2D grid, the implementation of the ReLU function remains constant in T -depth. For other activation functions, such as Sigmoid, Tanh, ELU, Swish, Softmax, and GELU, which are not conveniently implemented directly with arithmetic circuits, we employ quantum Lookup tables. We have balanced the number of input qubits against implementation precision and the trade-off between circuit T -depth and ancilla count to enhance the applicability of our results. Our method demonstrates advantages through a specific analysis of T -depth comparisons.

Looking ahead, we anticipate leveraging our activation function implementations in the implementation of quantum neural networks, thereby paving the way for future research directions in this domain.

ACKNOWLEDGMENTS

This work is supported in part by the National Research Foundation, Singapore, and A*STAR under its CQT Bridging Grant and its Quantum Engineering Programme under Grant No. NRF2021-QEP2-02-P05. This work is supported in part by an MoE Tier-1 grant. This work was supported in part by the National Natural Science Foundation of China Grants No. 62325210, as well as the Strategic Priority Research Program of the Chinese Academy of Sciences Grant No. XDB28000000. Additionally, this work was funded by the China Scholarship Council.

-
- [1] OpenAI *et al.*, GPT-4 technical report, [arXiv:2303.08774](https://arxiv.org/abs/2303.08774).
 - [2] C. Affonso, A. L. D. Rossi, F. H. A. Vieira, A. C. P. de Leon Ferreira *et al.*, Deep learning for biological image classification, *Expert Syst. Appl.* **85**, 114 (2017).
 - [3] D. Nie, R. Trullo, J. Lian, L. Wang, C. Petitjean, S. Ruan, Q. Wang, and D. Shen, Medical image synthesis with deep convolutional adversarial networks, *IEEE Trans. Biomed. Eng.* **65**, 2720 (2018).
 - [4] X. Yi and P. Babyn, Sharpness-aware low-dose CT denoising using conditional generative adversarial network, *J. Digital Imaging* **31**, 655 (2018).
 - [5] S. Li, W. Song, L. Fang, Y. Chen, P. Ghamisi, and J. A. Benediktsson, Deep learning for hyperspectral image classification: An overview, *IEEE Trans. Geosci. Remote Sens.* **57**, 6690 (2019).
 - [6] K. Gurney, *An Introduction to Neural Networks* (CRC Press, London, UK, 1997).
 - [7] M. Schuld, I. Sinayskiy, and F. Petruccione, An introduction to quantum machine learning, *Contemp. Phys.* **56**, 172 (2015).
 - [8] D. Peral-García, J. Cruz-Benito, and F. J. García-Peñalvo, Systematic literature review: Quantum machine learning and its applications, *Comput. Sci. Rev.* **51**, 100619 (2024).
 - [9] A. Abbas, D. Sutter, C. Zoufal, A. Lucchi, A. Figalli, and S. Woerner, The power of quantum neural networks, *Nat. Comput. Sci.* **1**, 403 (2021).
 - [10] P. Rebentrost, M. Mohseni, and S. Lloyd, Quantum support vector machine for big data classification, *Phys. Rev. Lett.* **113**, 130503 (2014).
 - [11] S. Lloyd, M. Mohseni, and P. Rebentrost, Quantum principal component analysis, *Nat. Phys.* **10**, 631 (2014).

- [12] J. Tilly, H. Chen, S. Cao, D. Picozzi, K. Setia, Y. Li, E. Grant, L. Wossnig, I. Rungger, G. H. Booth *et al.*, The variational quantum eigensolver: A review of methods and best practices, *Phys. Rep.* **986**, 1 (2022).
- [13] M. Benedetti, E. Lloyd, S. Sack, and M. Fiorentini, Parameterized quantum circuits as machine learning models, *Quantum Sci. Technol.* **4**, 043001 (2019).
- [14] L. Zhou, S.-T. Wang, S. Choi, H. Pichler, and M. D. Lukin, Quantum approximate optimization algorithm: Performance, mechanism, and implementation on near-term devices, *Phys. Rev. X* **10**, 021067 (2020).
- [15] M. Schuld, I. Sinayskiy, and F. Petruccione, The quest for a quantum neural network, *Quant. Inf. Process.* **13**, 2567 (2014).
- [16] K. H. Wan, O. Dahlsten, H. Kristjánsson, R. Gardner, and M. S. Kim, Quantum generalisation of feedforward neural networks, *npj Quantum Inf.* **3**, 36 (2017).
- [17] J. Zhao, Y.-H. Zhang, C.-P. Shao, Y.-C. Wu, G.-C. Guo, and G.-P. Guo, Building quantum neural networks based on a swap test, *Phys. Rev. A* **100**, 012334 (2019).
- [18] F. Tacchino, C. Macchiavello, D. Gerace, and D. Bajoni, An artificial neuron implemented on an actual quantum processor, *npj Quantum Inf.* **5**, 26 (2019).
- [19] I. Cong, S. Choi, and M. D. Lukin, Quantum convolutional neural networks, *Nat. Phys.* **15**, 1273 (2019).
- [20] M. Schuld, I. Sinayskiy, and F. Petruccione, Simulating a perceptron on a quantum computer, *Phys. Lett. A* **379**, 660 (2015).
- [21] F. M. de Paula Neto, T. B. Ludermir, W. R. de Oliveira, and A. J. da Silva, Implementing any nonlinear quantum neuron, *IEEE Trans. Neural Networks Learn. Syst.* **31**, 3741 (2019).
- [22] M. Maronese, C. Destri, and E. Prati, Quantum activation functions for quantum neural networks, *Quant. Inf. Process.* **21**, 128 (2022).
- [23] S. Yan, H. Qi, and W. Cui, Nonlinear quantum neuron: A fundamental building block for quantum neural networks, *Phys. Rev. A* **102**, 052421 (2020).
- [24] R. Livni, S. Shalev-Shwartz, and O. Shamir, On the computational efficiency of training neural networks, in *Advances in Neural Information Processing Systems 27: Annual Conference on Neural Information Processing Systems, Montreal, Quebec, Canada* (2014), pp. 855–863.
- [25] D. Gottesman, An introduction to quantum error correction and fault-tolerant quantum computation, in *Quantum Information Science and its Contributions to Mathematics, Proceedings of Symposia in Applied Mathematics* (American Mathematical Society, Providence, RI, 2010), Vol. 68, pp. 13–58.
- [26] S. Sharma, S. Sharma, and A. Athaiya, Activation functions in neural networks, *Int. J. Eng. Appl. Sci. Technol.* **4**, 310 (2020).
- [27] S. Goel and A. R. Klivans, Learning neural networks with two nonlinear layers in polynomial time, in *Proceedings of the Thirty-Second Conference on Learning Theory*, edited by A. Beygelzimer and D. Hsu (PMLR, USA, 2019), pp. 1470–1499.
- [28] D. Hendrycks and K. Gimpel, Gaussian error linear units (GELUs), [arXiv:1606.08415](https://arxiv.org/abs/1606.08415).
- [29] R. Krishnakumar, M. Soeken, M. Roetteler, and W. Zeng, A Q# implementation of a quantum lookup table for quantum arithmetic functions, in *Proceedings of the 2022 IEEE/ACM Third International Workshop on Quantum Computing Software (QCS)* (IEEE, USA, 2022), pp. 75–82.
- [30] S. Krinner, N. Lacroix, A. Remm, A. Di Paolo, E. Genois, C. Leroux, C. Hellings, S. Lazar, F. Swiadek, J. Herrmann *et al.*, Realizing repeated quantum error correction in a distance-three surface code, *Nature (London)* **605**, 669 (2022).
- [31] Google Quantum AI, Suppressing quantum errors by scaling a surface code logical qubit, *Nature (London)* **614**, 676 (2023).
- [32] D. Bluvstein, S. J. Evered, A. A. Geim, S. H. Li, H. Zhou, T. Manovitz, S. Ebadi, M. Cain, M. Kalinowski, D. Hangleiter *et al.*, Logical quantum processor based on reconfigurable atom arrays, *Nature (London)* **626**, 58 (2024).
- [33] <https://github.com/ziwei1998/quantum-circuit-for-activation-functions>.
- [34] Y. Takahashi and S. Tani, Power of uninitialized qubits in shallow quantum circuits, *Theor. Comput. Sci.* **851**, 129 (2021).
- [35] M. A. Nielsen and I. L. Chuang, *Quantum Computation and Quantum Information* (Cambridge University Press, Cambridge, UK, 2010).
- [36] P. Gokhale, S. Koretsky, S. Huang, S. Majumder, A. Drucker, K. R. Brown, and F. T. Chong, Quantum fan-out: Circuit optimizations and technology modeling, in *Proceedings of the 2021 IEEE International Conference on Quantum Computing and Engineering (QCE)* (IEEE, USA, 2021), pp. 276–290.
- [37] M. Fang, S. Fenner, F. Green, S. Homer, and Y. Zhang, Quantum lower bounds for fanout, *Quantum Inf. Comput.* **6**, 46 (2006).
- [38] T. H. Cormen, C. E. Leiserson, R. L. Rivest, and C. Stein, *Introduction to Algorithms* (MIT Press, Cambridge, MA, 2022).
- [39] IBM Quantum Platform (2024), <https://quantum.ibm.com/>.
- [40] Y. Wu, W.-S. Bao, S. Cao, F. Chen, M.-C. Chen, X. Chen, T.-H. Chung, H. Deng, Y. Du, D. Fan *et al.*, Strong quantum computational advantage using a superconducting quantum processor, *Phys. Rev. Lett.* **127**, 180501 (2021).
- [41] S. J. Evered, D. Bluvstein, M. Kalinowski, S. Ebadi, T. Manovitz, H. Zhou, S. H. Li, A. A. Geim, T. T. Wang, N. Maskara *et al.*, High-fidelity parallel entangling gates on a neutral-atom quantum computer, *Nature (London)* **622**, 268 (2023).
- [42] W. Kahan, IEEE Standard for Floating-Point Arithmetic, *IEEE Std 754-2019 (Revision of IEEE 754-2008)* (IEEE, 2019), pp. 1–84.
- [43] G. H. Low, V. Kliuchnikov, and L. Schaeffer, Trading T-gates for dirty qubits in state preparation and unitary synthesis, *Quantum* **8**, 1375 (2024).
- [44] A. Barenco, C. H. Bennett, R. Cleve, D. P. DiVincenzo, N. Margolus, P. Shor, T. Sleator, J. A. Smolin, and H. Weinfurter, Elementary gates for quantum computation, *Phys. Rev. A* **52**, 3457 (1995).
- [45] C. Gidney, Halving the cost of quantum addition, *Quantum* **2**, 74 (2018).
- [46] A. Gilyén, Y. Su, G. H. Low, and N. Wiebe, Quantum singular value transformation and beyond: Exponential improvements for quantum matrix arithmetics, in *Proceedings of the 51st Annual ACM SIGACT Symposium on Theory of Computing, Phoenix, AZ, USA* (Association for Computing Machinery, New York, 2019), pp. 193–204.