Fast High-Fidelity Readout of a Single Trapped-Ion Qubit via Machine-Learning Methods

Zi-Han Ding, Jin-Ming Cui,* Yun-Feng Huang, Chuan-Feng Li,† Tao Tu, and Guang-Can Guo
CAS Key Laboratory of Quantum Information, University of Science and Technology of China, Hefei 230026, China and CAS Center For Excellence in Quantum Information and Quantum Physics, University of Science and Technology of China, Hefei 230026, People's Republic of China

(Received 25 October 2018; revised manuscript received 19 May 2019; published 22 July 2019)

In this work, we introduce machine-learning (ML) methods to implement fast and high-fidelity readout of a trapped-ion qubit on a hardware module, which is based on field-programmable gate arrays (FPGAs) and an ARM (Advanced RISC Machines) processor. An average readout fidelity of 99.5% (with 10^5 magnitude trials) within 171 μs is achieved in experiments. Different ML architectures including convolutional neural networks and fully connected neural networks are implemented to compare with traditional methods, demonstrating higher fidelity and more robust readout results in a relatively short time with the proposed methods. Our hardware implementation of the proposed ML methods with FPGAs improves the readout efficiency to a higher level, through reducing the communication time between the trapped-ion system and central processing units (CPUs) in general personal computer (PCs) and the hardware can be devolved to a functional module, which is compatible with the real-time readout and feedback control of qubit states.

DOI: 10.1103/PhysRevApplied.12.014038

I. INTRODUCTION

Quantum computing [1,2] requires high-quality quantum-logic gates, including highly accurate qubit operations and readout. The fast and high-fidelity readout [3] of quantum states [4] is essential for fault-tolerant quantum computation. The state (bright or dark) of a trapped ion is derived through the qubit readout process.

Traditional methods for qubit-state discrimination, such as the threshold method, the maximum-likelihood method [5], or the adaptive-maximum-likelihood method [6], cannot guarantee high-accuracy readout within a short detection time. The threshold method simply discriminates a qubit state with the sum of the photon-counting sequence, regardless of the inner patterns of the data sequences. However, the state flip caused by long-time laser interaction during state detection, which makes the bright state of $^{171}Yb^{+}$ "jump" to the dark state, is hardly recognized by the threshold method. The maximum-likelihood method has the ability to recognize state flips during the detection period and therefore guarantees higher discrimination accuracy. However, the detection time consumption for the statistical process of the maximum-likelihood method is large. Furthermore, the threshold and maximum-likelihood methods grasp relatively few inner patterns or little of the information contained in photon-counting sequences, which makes it hard for them to guarantee high accuracy within a short detection time.

Using recently popular machine-learning (ML) tools [7], we design methods to guarantee both low time consumption and high-fidelity readout on ion-trap systems. Experiments have shown that ML methods such as neural networks are good at capturing features and patterns in a sequence of data. We experimentally test our proposed ML-assisted single-qubit readout methods on a trapped $^{171}Yb^{+}$ ion and obtain faster and higher-fidelity readout than traditional methods.

Real-time state discrimination is important in applications such as fast feedback control of quantum-state operations. Considering that the field-programmable-gate-array- (FPGA) [8] based system has the ability to process data faster [9], we implement real-time state discrimination with ML algorithms on an embedded-hardware system instead of the traditional CPU or graphics processing unit (GPU) implementation [10]. We apply FPGA-based preprocessing and ARM-processor-based feedforward neural networks for fast qubit readout. The time properties and the fidelity are experimentally measured on our embedded-hardware qubit readout system. In recent research, the ML method has also been applied in multi-ion state readout [11], to reduce the cross-talk error between different qubits. Algorithms including a support-vector machine are applied in the classification of quantum-measurement trajectories for the improvement of qubit measurements in superconducting systems [12]. However, our work focuses

*jmcui@ustc.edu.cn
†cfli@ustc.edu.cn

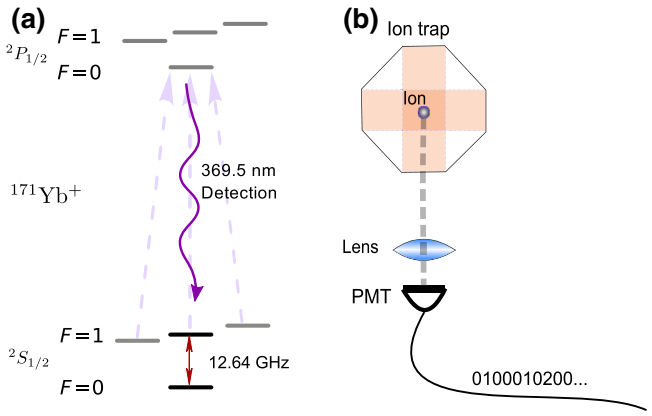

FIG. 1. The energy levels and single-qubit readout on the $^{171}Yb^+$ trapped-ion system. A laser of 369.53 nm is used to excite the bright state $^2S_{1/2}\,|F=1\rangle$, so as to spontaneously radiate fluorescence, which can be collected via a NA = 0.4 lens and detected by the following photomultiplier tube (PMT). The output of the PMT is processed to be the photon-counting numbers for the purpose of ion-state (also called a qubit) discrimination.

on leveraging the sequential information in single-qubit detection, instead of the cross talk of the multi-ion state. Also, with our hardware implementation using FPGAs and ARM processors, the time performance is improved so as to be compatible with real-time feedback control of ion states.

II. ML METHODS FOR SINGLE-QUBIT READOUT IN AN ION TRAP

The energy levels that we use in a $^{171}Yb^+$ radio-frequency Paul trap [13] are $^2S_{1/2}\,|F=0\rangle$ as the dark state and $^2S_{1/2}\,|F=1\rangle$ as the bright state, as shown in Fig. 1(a). The magnetic dipole transition between these two energy levels corresponds to 12.6-GHz microwave [14] operation. The 369.53-nm detection laser can arouse resonance between two states: $^2S_{1/2}\,|F=1\rangle \leftrightarrow {}^2P_{1/2}\,|F=0\rangle$. A state flip could be caused by a microwave $\pi/2$ pulse. The excited state $^2P_{1/2}\,|F=0\rangle$ will sustain a spontaneous radiation process to emit fluorescence, which can be detected by the photomultiplier tubes (PMT). However, the state $^2S_{1/2}\,|F=0\rangle$ cannot be excited because of the 12.6-GHz detuning. The readout process of a single qubit is shown in Fig. 1(b). Experimentally, we use an objective lens with NA = 0.4 for the detection of photons.

Visualized results of the single-qubit readout are shown in Fig. 2. We experimentally measure the dark state and the bright state ten times each with the PMT, to obtain a 100-sub-bin (3 μs per sub-bin) data sequence each time. These sequences are the raw data used to determine whether it is a dark- or bright-state qubit.

ML methods have been proved to have notable effects on applications such as sequence data classification, especially with supervised learning. Considering the low cost of generating labeled data, we test several supervised-learning methods, including fully connected neural networks (NNs), convolutional neural networks (CNNs) [15], recurrent neural networks (RNNs) [16], support-vector machines (SVMs) [17], logistic regression [18], and decision-tree classifiers [19] to compare with traditional methods such as the threshold and maximum-likelihood methods [6]. As the inputs of supervised learning are data organized into time series, we can naturally apply RNN-based methods. However, the RNN consumes too much time, as shown in the comparisons. We could neglect the time properties of the input data and apply a fully connected NN to conduct the classification, or even treat it as "spatial patterns" with a CNN, which shows better effects. Furthermore, SVM or regression-based methods could also be applied for classifying the high-dimensional input data.

Figure 3 shows the architecture of the one-dimensional (1D) CNN that we apply for the experiments, similar to LENET-5 [20]. The structure that we apply is fine tuned with considerations regarding both accuracy and computation time. It has two convolutional layers, two pooling layers,

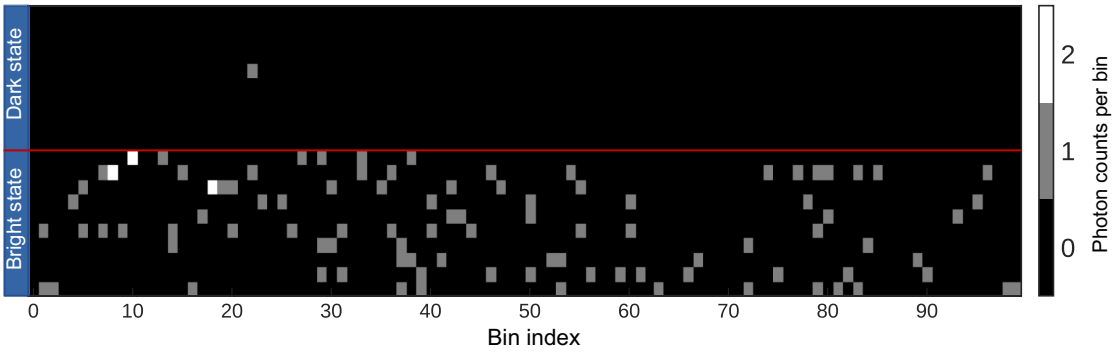

FIG. 2. The single-qubit dark-state and bright-state readout on the $^{171}Yb^+$ trapped-ion system. The diagram shows the readout experiment, carried out ten times, for both dark-state (0–9) and bright-state (10–19) sequences, each with photon counts in 100 sub-bins.

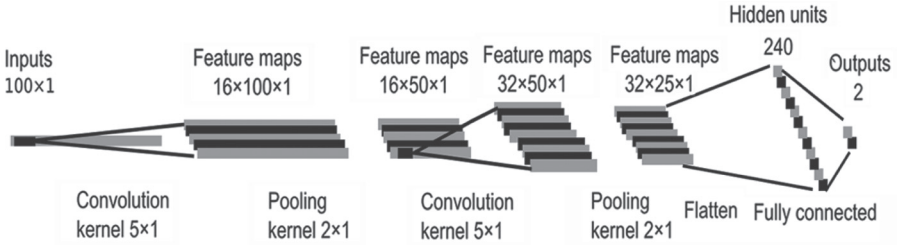

FIG. 3. The architecture of the CNN for single-qubit readout in a test. For the two units of the output layer, y_1 and y_2, if $y_1 > y_2$, the state is inferred to be a bright state; otherwise, it is a dark state.

and two fully connected layers. The dimension of the input is 100×1 (downsampling into 50×1, 25×1 through max-pooling of size 2×1 after each convolution). The first convolutional layer has 16 feature maps of size 100×1, with a convolution kernel of 5×1. The second convolutional layer has 32 feature maps of size 50×1. The neural units in the fully connected layer and the output layer are of size 240 and size 2, respectively. The size of the minibatch in the gradient descent process is 100. The ADAMOPTIMIZER [21] is applied with a decayed learning rate of $1 \times 10^{-3} \sim 1 \times 10^{-4}$, so that the model can be trained to reach its optimum within2×10^5 samples. The loss function that we apply in our model is the mean-square error (MSE) between the output inferences and the target labels (with a similar learning effect as cross-entropy loss in our model), as follows:

$$\text{loss} = \frac{1}{N}\sum_{x\in X}(y_x - \hat{y_x})^2, \tag{1}$$

where X is the set of input photon-counting sequences and N is its cardinality, y_x are the output inference values of the neural network, and $\hat{y_x}$ is the true label of the corresponding input sequence.

The structures of the other ML methods are not further described here, but they are all experimentally optimized. The inputs of the ML methods are the original photon-counting sequences, which are 100-sub-bin counts used to be detect for each state. The training data of 2×10^5 samples is generated directly on the experimental systems. We experimentally prepare the bright and dark states with the $\pi/2$ pulse alternately on and off, which helps to eliminate systematic errors. Each sample is a binary sequence of 300 μs measured from the PMT for single state preparation, to be divided into different numbers of sub-bins on demand (e.g., 100 sub-bins). The labels are set alternately on samples to be consistent with the experimental preparations. With fast preparation of bright and dark states in ion traps, we can obtain a large batch of training data in a short period of time. After training the CNN model with experimental data, we subsequently use the model to classify the states generated in experiments in the same way as above. The outputs of the CNN are predictions of bright or dark states. Therefore, the CNN is an end-to-end model for qubit discrimination, as well as the other ML methods.

We experimentally test our proposed ML methods and traditional methods using the same data set. As shown in Table I, among all the ML and traditional methods, the CNN method achieves the highest rate of accuracy and relatively robust performance with the input of 100 sub-bins of photon-counting data. To investigate the time properties of different methods, a test with a data set of 2×10^5 samples was undertaken for each method with cross validation. The time consumption of CPU computation for qubit discrimination is also shown in Table I. The faster the inference process runs, the sooner we can obtain the states of the ion. Traditional methods such as maximum likelihood need a long computation time for complete numerical statistics of photon bin counts, in order to maintain high accuracy. However, most ML methods demonstrate a lower demand on computation time, except for RNN methods. With a significantly higher accuracy of qubit discrimination, CNN and fully connected neural-network methods only need relatively short computation times, which makes them promising for fast readout and feedback control of single qubit. We will further speed up the neural-network method with hardware implementation in Sec. III.

It is worth noting that the neural-network methods have a higher accuracy than the maximum-likelihood method. This could be explained by the fact that the formula applied in traditional maximum-likelihood methods in Ref. [6] is an approximated version of the true formula of the probability of a dark state, derived from the master equation as in Ref. [5]. However, it is too complicated to directly apply the true formula for estimation of the dark-state probability with many more variables. Therefore,

TABLE I. A comparison of the accuracy and inference time of different single-qubit discrimination methods.

Method	Accuracy (%)	Time (s)
Threshold	99.248 ± 0.07	1.563
Maximum likelihood	99.311 ± 0.12	4678.839
Fully connected NN	99.411 ± 0.13	2.865
CNN	$\mathbf{99.413 \pm 0.10}$	7.065
RNN	99.364 ± 0.10	914.113
SVM	99.341 ± 0.10	4.282
Logistic regression	99.123 ± 0.11	4.137
Decision-tree classifier	98.010 ± 0.30	1.504

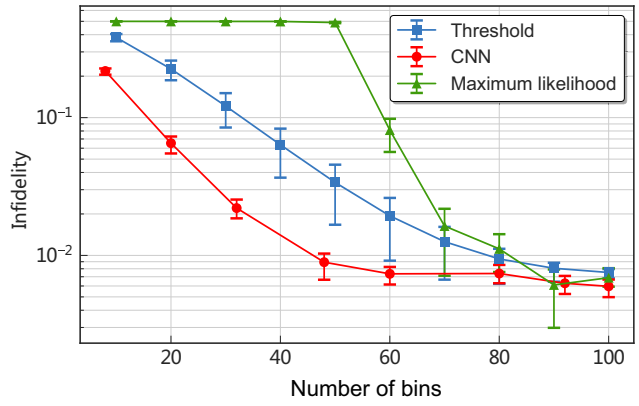

FIG. 4. The infidelity of single-qubit readout with different numbers of sub-bins in a test. The proposed CNN method can achieve 99% accuracy within 43 sub-bins (approximately 129 μs), while the threshold and maximum-likelihood methods need more than 80 sub-bins. And the CNN method has the more robust readout performance.

the maximum-likelihood method in qubit discrimination is inaccurate to some extent, while the proposed ML (neural-network) methods have a great potential to eliminate these errors, with a sufficient capability to approximate any continuous functions [22], as long as there are enough degrees of freedom in the layers and nodes with proper activation functions, which guarantees their higher accuracy compared to conventional methods such as maximum likelihood.

Figure 4 shows the relationship between the infidelity =1-fidelity, where fidelity is a value of [0,1] and the number of sub-bins (fixed sub-bin time) for the CNN and two traditional methods. To achieve 99% accuracy, only 43 sub-bins (approximately 129 μs) are needed by the CNN method, while the threshold method needs 80 sub-bins (approximately 240 μs). The proposed CNN method only needs half the amount of data needed by the traditional methods to guarantee the same accuracy, which significantly reduces the detection time needed for qubit discrimination. With fewer bin data, ML methods could make the qubit-discrimination process even faster. Additionally, the error bars in the diagram confirm the more robust performance of the CNN compared with the other two methods, which is essential for high-fidelity one-shot single-qubit readout.

III. EMBEDDED-QUBIT READOUT SYSTEM

To discriminate the state in real time with a single-shot readout, we need to record the photon counts and conduct the feedforward neural-network algorithm in each detection. Special hardware is needed to support this feature [23]. The overall computation process is embedded into the chip Zynq-AX7020, which has an ARM (Cortex-A9, 767 MHz) processor and a FPGA (XC7Z020-2CLG400I). The transistor-transistor logic (TTL) signals from the PMT are directly sent to the chip. Two main functions are accomplished on the chip: (1) signal transformation from the TTL to the digital photon counts and its storage in registers; and (2) feedforward neural-network implementation on an ARM processor with FPGA computation acceleration. With the above settings, the average processing time for single-sample discrimination is significantly decreased.

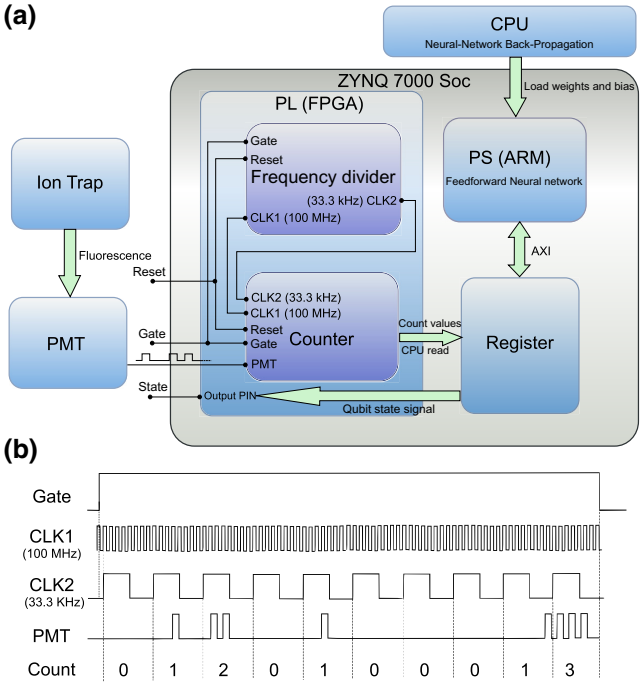

FIG. 5. The fast qubit readout system as used in the experiment. (a) The scheme of the overall system. The inputs of our proposed embedded-qubit readout system are TTL signals from the PMT, which are acquired through detection of the bright-state fluorescence in the ion-trap system. The gate signal is an enable-disable signal to control the counting process. A frequency divider and a counter are implemented on a FPGA and the implementation of a feedforward neural network is programmed into the ARM processor. Every time the counter counts a photon number in a sub-bin time, the PL triggers an interrupt to stop the loop of the software program in the PS and sends the number through the registers to the PS. The parameters for the weights and the bias of the neural network are pretrained on a CPU or GPU and loaded on the Zynq development board. Whenever the PS receives a signal that denotes the end of a bin time, it automatically starts the inference process and outputs the state of the ion to the pin. (b) The waveform-sequence diagram of the controlling signals in an experiment. The 33.3-kHz clock is generated from the 100-MHz system clock and serves as the controlling signals of the counter. The extraneous 1.67-kHz gate signal sets the time window (300 μs high-level voltage for the detection time) of the single-sample photon counting and it is the controlling signals of both the counter and the frequency divider. Each rising edge of the PMT signals is detected to be one photon count within a one-sub-bin time 30 μs, which is in accordance with the period of the generated 33.3-kHz clock.

The diagram of the embedded system on a Zynq-7000 development board for fast qubit readout is shown in Fig. 5(a). A frequency divider and an analog-to-digital counter for the TTL signals are implemented on a FPGA through hardware programming. Experimentally, we set the sub-bin time to be 30 μs, instead of the 3 μs used in the above simulation tests, to achieve faster qubit discrimination. Correspondingly, the number of input units becomes ten instead of 100, to keep the total bin time unchanged. Our tests show that, to some extent, the merging of inputs does not significantly decrease the effects of neural networks for qubit discrimination. The system clock for the programmable-logic unit (PL) is 100 MHz (approximately 10 ns). The frequency divider transforms this clock to a value of around 33.3 kHz, which is used as a sampling clock to record the photon counts within a sub-bin time of 30 μs. As shown in Fig. 5(b), the 33.3-kHz clock is generated from the 100-MHz system clock on board through the frequency divider, starting from the positive edge of the gate controlling the signal and ending with the negative edge of the gate signal. The maximum asynchronous error of the generated clock and the gate signal could be no more than 10 ns, which is the period of the 100-MHz system clock. The generated clock, together with the gate signal, controls the time pin on board of the counter. The falling edge of the gate signal will trigger an interrupt for the ARM processor to start the inference process and output the state of the ion.

The feedforward implementation of the fully connected neural network with two layers, of size 20 and size 2, respectively, and with a rectified linear unit (ReLU) [24] as an activation function, is set on the processing-system core unit (PS). $\mathrm{ReLU}(x) = x$ when input $x > 0$; otherwise, $\mathrm{ReLU}(x) = 0$. The weights and bias parameters are pretrained through the back-propagation [25] process on the CPU with a neural network exactly the same as the feedforward implementation on the PS. The ADAMOPTIMIZER is used for training this fully connected neural network with the decayed learning rate of $1 \times 10^{-3} \sim 1 \times 10^{-4}$. After it is trained to be optimal with about 2×10^5 samples, the weights and bias parameters are stored in files, to be loaded onto the development board for the inference process.

IV. EXPERIMENTAL RESULTS

We experimentally test the methods above, using an oscilloscope to display the detection results of the qubit readout as shown in Fig. 6. A bright-state and dark-state cross-fade sequence is prepared with a microwave $\pi/2$ pulse in the single-trapped-ion system, as shown in Fig. 6(a). The detection of the bright state is set to output a spike pulse signal after five sub-bin photon counts and a 21 μs inference time. The signal is actually a triple 1 μs pulse following the high-level gate signals, seen more clearly in Fig. 6(b). The enlarged images of readout results of

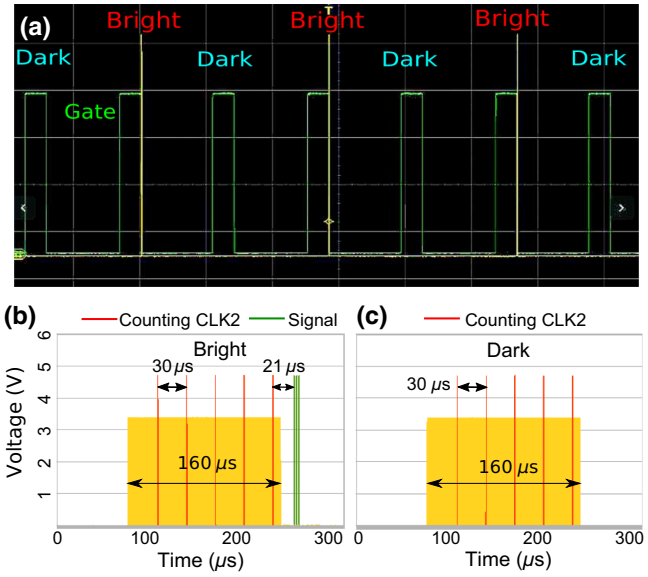

FIG. 6. Detection signals measured by an oscilloscope. (a) The experimental readout of a cross-fade bright and/or dark state sequence, with seven states in scope and five detection sub-bins for each state. (b),(c) Enlarged images of signals of a dark state and a bright state, respectively. The yellow rectangular wave represents the high-level gate signals of 160 μs (>30 × 5, to ensure five detection bins) and the red 1 μs pulse inside the gate signal is the sampling clock sign for each sub-bin. The green triple pulse right after the high-level gate signal in (b) is the signal for a bright state; its absence indicates a dark state.

the single-sample qubit readouts for the bright and dark states are shown in Figs. 6(b) and 6(c), with extra signals for each counting period (five red pulses). We could derive that the inference time for each sample is around 21 μs in the embedded system, as measured in Fig. 6(b). However, the average inference time with a CPU and a GPU on a PC is given as 72 μs in Table II, not to mention the considerable time consumption on communication between the PMT counter and the computer. Therefore, the embedded-hardware implementation of the fully connected neural-network method has a speed-up by a factor of more than 3 in comparison to general systems. The total time consumption for the single-sample qubit readout is about 171 μs (5 bins × 30 μs/bin + 21 μs = 171 μs).

TABLE II. The time properties of different computation architectures for single-sample qubit discrimination.

	Computation time (ms)	
Architecture	Loading files	Inference
CPU + GPU (test)	0.42	0.072
ARM (Test)	91.2	2.7
ARM + FPGA (test)	2.2	1.1
ARM + FPGA (experiment)	–	**0.021**

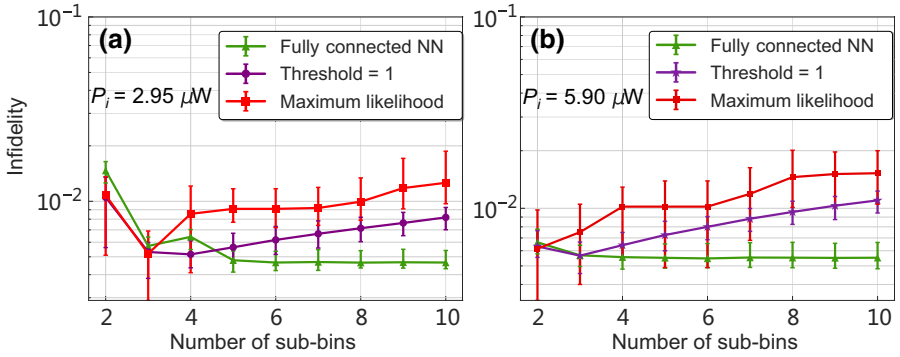

FIG. 7. A comparison of the different methods for single-qubit readout with incident laser powers of 2.95 μW and 5.90 μW.

As there is a saturation power for detecting the ion fluorescence, we might choose the optimum excitation laser power for detection. The detected fluorescence photon-counting rate n with excitation power P is as follows:

$$n(P) = n_0 \frac{x}{1+x} = n_0 \frac{P/P_0}{1+P/P_0},$$

where P_0 is the saturation power and n_0 denotes the saturation photon-counting rate. In our setup, P_0 is measured as 2.91 μW and n_0 is 1.39×10^5 counts/s. Comparisons of the onboard fully connected neural-network method and the threshold and maximum-likelihood methods for single-qubit readout are shown in Fig. 7, with different numbers of sub-bins and incident laser powers of 2.95 μW and 5.90 μW, respectively. The best accuracy of 99.53% is achieved with the incident laser power of 2.95 μW. Moreover, a faster qubit readout with only the data from two sub-bins as input could be achieved using a stronger laser of 5.90 μW, as well as maintaining a fidelity higher than 99%. The rise of the infidelity for the threshold method with more sub-bins shows that the dark state has a greater chance of being excited into a bright state with larger incident laser powers such as 2.95 μW and 5.90 μW, rather than using 1.26 μW as usual. More importantly, the robust performance of the neural-network method proves that it could capture these states' transformation features, which is reflected by the photon-counting sequence, in the process of state inference. Considering the time properties, fidelity, and robustness, the best performance is achieved by having fully connected neural networks on board with the data from five sub-bins (30 μs for each bin) as the input variables and a 2.95 μW 369.53 nm detecting laser in the system.

The time properties of different computation architectures for the single-qubit discrimination of one sample are shown in Table II. All of the architectures have realized the function of a feedforward fully connected neural network with ten input units, as described above. As shown in Fig. 5, the files containing the weights and bias of the neural network are loaded into the computation modules, including those for the GPU (test) and the ARM processor (test or experiment). Then the computation of the neural-network inference is accomplished on these architectures. The three tests are implemented using the data files of the digital photon-counting sequences, derived through measurements of the ion state before the tests; the experiment with an ARM processor and a FPGA on the Zynq-7000 development board is implemented with the PL part for processing raw inputs of PMT signals from the physical system, which are analog rather than digital. In other words, in the tests, ARM + FPGA only uses the PS part of the system in Fig. 5, while in the experiments, ARM + FPGA uses both the PL and the PS parts; and in the tests, the FPGA is for acceleration of the PS computation process, unlike the experiments.

V. CONCLUSION

To conclude, for single-qubit readout in a trapped-ion system, we can achieve 99.5% fidelity within around 171 μs per sample, using a FPGA- and ARM-processor-based fully connected neural-network method. The proposed ML methods for single-qubit readout reduce half of the inference time in simulation compared with conventional methods and the hardware implementation further improves the efficiency by a factor of at least 3, without even considering the longer communication time between the trapped-ion system and a PC with a general CPU and/or GPU. Moreover, higher fidelity and more robust performance could be achieved with convolutional neural networks and more sub-bins. On the other hand, the performance of the proposed ML methods is limited to some extent by the accuracy of our preparations of qubit states in the experiments. We could expect higher accuracy with more sophisticated and precise state preparation. In general, the ML methods are applied to grasp the more detailed features of the photon-counting sequences for a higher discrimination accuracy with a shorter detection time; while the FPGA- and ARM-processor-based hardware provides a faster process of neural-network inference. The good time performance of the proposed framework enables almost real-time readout and feedback control of the qubit state, which is of great significance for more advanced quantum

computation beyond the single qubit. Furthermore, our embedded-qubit readout method could be extended to synchronous multi-ion-state readout and fast feedback control of qubit states in trapped-ion systems, which is promising for more flexible quantum-gate operations in the future.

ACKNOWLEDGMENTS

We would like to thank Yao Teng, Chang-Qing Feng, and Yan-Wen Wu for their experimental support of the FPGA board development and for helpful discussions. This work is supported by the National Key Research and Development Program of China (Grants No. 2017YFA0304100 and No. 2016YFA0302700), the National Natural Science Foundation of China (Grants No. 61327901, No. 11774335, No. 11474268, No. 11474270, No. 11734015, and No. 11821404), the Key Research Program of Frontier Sciences, CAS (Grant No. QYZDY-SSW-SLH003), the Fundamental Research Funds for the Central Universities (Grants No. WK2470000026, No. WK2470000027, and No. WK2470000028), the Anhui Initiative in Quantum Information Technologies (Grants No. AHY020100 and No. AHY070000), and the National Program for Support of Topnotch Young Professionals (Grant No. BB2470000005).

[1] M. A. Nielsen and I. Chuang, *Quantum Computation and Quantum Information* (AAPT, 2002).
[2] J. I. Cirac and P. Zoller, Quantum Computations with Cold Trapped Ions, Phys. Rev. Lett. **74**, 4091 (1995).
[3] T. P. Harty, D. T. C. Allcock, C. J. Ballance, L. Guidoni, H. A. Janacek, N. M. Linke, D. N. Stacey, and D. M. Lucas, High-Fidelity Preparation, Gates, Memory, and Readout of a Trapped-Ion Quantum Bit, Phys. Rev. Lett. **113**, 220501 (2014).
[4] J. Kelly, R. Barends, A. Fowler, A. Megrant, E. Jeffrey, T. White, D. Sank, J. Mutus, B. Campbell, Y. Chen, *et al.*, State preservation by repetitive error detection in a superconducting quantum circuit, Nature **519**, 66 (2015).
[5] C. E. Langer, Ph.D. thesis, University of Colorado (2006).
[6] A. H. Myerson, D. J. Szwer, S. C. Webster, D. T. C. Allcock, M. J. Curtis, G. Imreh, J. A. Sherman, D. N. Stacey, A. M. Steane, and D. M. Lucas, High-Fidelity Readout of Trapped-Ion Qubits, Phys. Rev. Lett. **100**, 200502 (2008).
[7] Y. LeCun, Y. Bengio, and G. Hinton, Deep learning, Nature **521**, 436 (2015).
[8] S. D. Brown, R. J. Francis, J. Rose, and Z. G. Vranesic, *Field-Programmable Gate Arrays* (Springer Science & Business Media, 2012), Vol. 180.
[9] C. Yang, C. Feng, W. Dong, D. Jiang, Z. Shen, S. Liu, and Q. An, Alpha-gamma discrimination in BaF_2 using FPGA-based feedforward neural network, IEEE Trans. Nucl. Sci. **64**, 1350 (2017).
[10] J. Misra and I. Saha, Artificial neural networks in hardware: A survey of two decades of progress, Neurocomputing **74**, 239 (2010).
[11] A. Seif, K. A. Landsman, N. M. Linke, C. Figgatt, C. Monroe, and M. Hafezi, Machine learning assisted readout of trapped-ion qubits, J. Phys. B: At., Mol. Opt. Phys. **51**, 174006 (2018).
[12] E. Magesan, J. M. Gambetta, A. D. Córcoles, and J. M. Chow, Machine Learning for Discriminating Quantum Measurement Trajectories and Improving Readout, Phys. Rev. Lett. **114**, 200501 (2015).
[13] S. Debnath, N. M. Linke, C. Figgatt, K. A. Landsman, K. Wright, and C. Monroe, Demonstration of a small programmable quantum computer with atomic qubits, Nature **536**, 63 (2016).
[14] C. Ospelkaus, U. Warring, Y. Colombe, K. Brown, J. Amini, D. Leibfried, and D. Wineland, Microwave quantum logic gates for trapped ions, Nature **476**, 181 (2011).
[15] A. Krizhevsky, I. Sutskever, and G. E. Hinton, Imagenet classification with deep convolutional neural networks, in *Advances in Neural Information Processing Systems* (2012), p. 1097.
[16] S. Hochreiter and J. Schmidhuber, Long short-term memory, Neural Comput. **9**, 1735 (1997).
[17] M. A. Hearst, S. T. Dumais, E. Osuna, J. Platt, and B. Scholkopf, Support vector machines, IEEE Intell. Syst. Appl. **13**, 18 (1998).
[18] D. W. Hosmer Jr, S. Lemeshow, and R. X. Sturdivant, *Applied Logistic Regression* (John Wiley & Sons, 2013), Vol. 398.
[19] S. R. Safavian and D. Landgrebe, A survey of decision tree classifier methodology, IEEE Trans. Syst., Man, Cybern. **21**, 660 (1991).
[20] Y. LeCun, L. Bottou, Y. Bengio, and P. Haffner, Gradient-based learning applied to document recognition, Proc. IEEE **86**, 2278 (1998).
[21] D. P. Kingma and J. Ba, Adam: A method for stochastic optimization, arXiv:1412.6980 (2014).
[22] B. C. Csáji, Master's thesis, Eötvös Loránd University (2001).
[23] Program available on GitHub: https://github.com/quantumiracle/On˙board˙FNN˙qubit˙discrimination.
[24] V. Nair and G. E. Hinton, Rectified linear units improve restricted boltzmann machines, in *Proceedings of the 27th International Conference on Machine Learning (ICML-10)* (2010), p. 807.
[25] Y. LeCun, D. Touresky, G. Hinton, and T. Sejnowski, in *Proceedings of the 1988 Connectionist Models Summer School* (CMU, Morgan Kaufmann, Pittsburgh, PA, 1988), Vol. 1, p. 21.