Machine learning assisted readout of trapped-ion qubits

Alireza Seif[1] , **Kevin A Landsman**[1,2], **Norbert M Linke**[1,2],
Caroline Figgatt[1,2], **C Monroe**[1,2,3] **and Mohammad Hafezi**[1,4]

[1] Joint Quantum Institute and Department of Physics, University of Maryland, College Park, MD 20742, United States of America
[2] Joint Center for Quantum Information and Computer Science, University of Maryland, College Park, MD 20742, United States of America
[3] IonQ, Inc., College Park, MD 20740, United States of America
[4] Department of Electrical and Computer Engineering and Institute for Research in Electronics and Applied Physics, University of Maryland, College Park, MD 20742, United States of America

E-mail: seif@umd.edu

Abstract

We reduce measurement errors in a quantum computer using machine learning techniques. We exploit a simple yet versatile neural network to classify multi-qubit quantum states, which is trained using experimental data. This flexible approach allows the incorporation of any number of features of the data with minimal modifications to the underlying network architecture. We experimentally illustrate this approach in the readout of trapped-ion qubits using additional spatial and temporal features in the data. Using this neural network classifier, we efficiently treat qubit readout crosstalk, resulting in a 30% improvement in detection error over the conventional threshold method. Our approach does not depend on the specific details of the system and can be readily generalized to other quantum computing platforms.

Supplementary material for this article is available online

(Some figures may appear in colour only in the online journal)

Quantum computing tasks involve quantum state preparation, time evolution, and measurement, accompanied by errors in all the three stages. To detect and correct errors during the evolution, quantum error correcting codes are used [1–3]. These codes rely on redundant encoding of quantum information, which makes it possible to measure syndromes and fix errors. Measurement errors not only affect the outcome of the computation, but they also limit the task of error correction. Consequently, in addition to high-fidelity operations, high quality multi-qubit readout is essential for realizing a fault tolerant quantum computer.

The quantum measurement process always involves the interaction with an external classical system. For example, collecting fluorescence from a trapped-ion in a cycling transition can determine the state of the qubit [4]. In superconducting qubits, a probe signal is injected to the system through a resonator, and the phase of the output signal is used to infer the state of the qubit [5]. Spontaneous decay and excitation during the external probe can be major sources of qubit measurement errors [6]. When scaling up, the measurement signal from a qubit can be altered by the state of other qubits through crosstalk. To address this issue, one can assume an error model and infer the correct qubit states by using statistical properties of the measured data [7].

Machine learning (ML) [8] techniques have recently become popular tools for exploring physical phenomena. For example, artificial neural networks [9] are now a powerful method for simulating the dynamics of many-body quantum systems [10, 11]. These neural networks can efficiently represent a wide class of highly correlated states [12–14], and can facilitate quantum state tomography [15, 16]. They are also used to detect errors and decode quantum error correcting codes [17–20], and to classify phases of matter [21–23]. In addition to
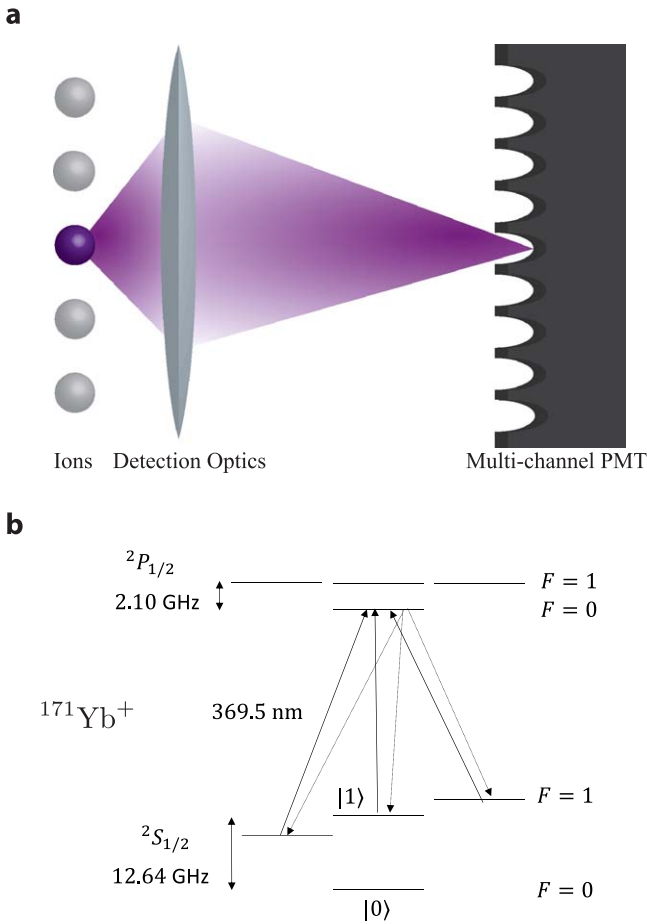
a

Ions   Detection Optics   Multi-channel PMT

b

Figure 1. The setup and readout scheme for the trapped-ion quantum computer. (a) Schematic of our experimental setup. A single ion fluoresces inside an ion trap and its radiated photons are collected by a 0.37 NA lens. This fluorescence is then imaged onto a single channel of a multi-channel photo-multiplier tube detector. (b) Energy levels in the ^{171}Yb$^+$ atomic system used for fluorescence detection. If the qubit state is $|0\rangle$, the applied detection laser is off resonance and nearly no photons are scattered. If the state is $|1\rangle$ the transition is on resonance and the ion fluoresces strongly.

neural networks, other ML methods, such as principal component analysis and clustering, have been used for various tasks from classifying phases of matter [24] to discriminating measurement trajectories for improved single qubit readout [25].

In this work, we exploit the versatility of ML techniques to increase the fidelity of multi-qubit measurements. While the problem of crosstalk can be partially addressed by careful statistical analysis of the data, it requires certain assumptions about the error model, which makes the integration of additional spatial and temporal features difficult. In our approach, the machine is 'trained' to infer the states from the measurement results without prior knowledge of the error model. This ML method can therefore be readily generalized to other quantum computing platforms.

We study the detection accuracy of a chain of ^{171}Yb$^+$ ions confined in an rf Paul trap [26]. The qubit is defined by the hyperfine-split ground states of the $^2S_{1/2}$ manifold: $|0\rangle = |F = 0, m_F = 0\rangle$ and $|1\rangle = |F = 1, m_F = 0\rangle$ (see figure 1). Furthermore, we can take advantage of the $^2P_{1/2}$
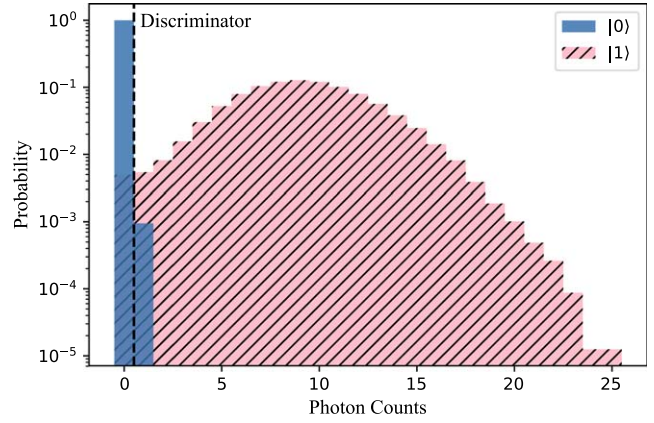

Figure 2. The histogram of observed photons for an integration time of 150 μs for state detection. The photon counts follow a near-Poissonian distribution, in which the state $|0\rangle$ (solid blue) gives a mean close to zero, while the state $|1\rangle$ (shaded pink) results in nine photons on average.

level to accomplish both state preparation and measurement (SPAM) with high-fidelity.

Qubit initialization is achieved by optical pumping via the $^2P_{1/2}$, $F = 1$ manifold. The qubit readout, on the other hand, is performed by state-dependent fluorescence detection [6] (see figure 2). Specifically, we apply a laser beam resonant with the $^2S_{1/2}$, $F = 1 \rightarrow {}^2P_{1/2}$, $F = 0$ cycling transition, and collect ion fluorescence. While the beam is on, a qubit in $|1\rangle$ will scatter photons. In contrast, a qubit in $|0\rangle$ remains dark since the light is 14.7 GHz detuned from the nearest transition with a natural linewidth of about 20 MHz. The ion fluorescence is collected by a 0.37 NA lens and each ion in the chain is imaged onto a separate channel of a 32-channel photo-multiplier tube (PMT) [26].

The histogram of the photon counts in some integration time follows a near-Poissonian distribution, centered around 0 for state $|0\rangle$ (the 'dark' state) and nine counts for state $|1\rangle$ (the 'bright' state) following a 150 μs integration time. The deviations from Poissonian statistics indicate the error mechanisms in this readout scheme. The dark state histogram includes a small contribution at higher counts due to off-resonant dark-to-bright pumping during the detection step [27]. More importantly, the bright-state histogram has a non-Poissonian tail towards lower photon counts due to off-resonant excitation to the $^2P_{1/2}$, $F = 1$ manifold, detuned by 2.1 GHz, from which decay to $|0\rangle$ is possible [27]. By choosing an optimal collection time, 150 μs in our system, the overlap between the photon distributions corresponding to $|0\rangle$, and $|1\rangle$ can be minimized. Thus, by discriminating the two distributions one can deduce the state of the qubit. One of the commonly used techniques to distinguish between these distributions is a simple threshold discriminator, where instances with observed photon counts greater than the threshold are taken to be $|1\rangle$, and those below to be $|0\rangle$ (see figure 2). This method works very well in the single qubit case and results in a detection fidelity, that is $\mathcal{F} = p$(measured x|prepared x), of 99.4% for $|1\rangle$, and 99.6% for $|0\rangle$, which gives an average detection fidelity of 99.5% in our setup. This error can be reduced by increasing the collection angle of the objective
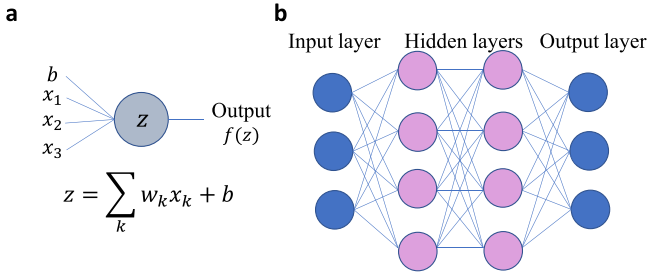
a

$$b$$
$$x_1$$
$$x_2$$
$$x_3$$

Output
$f(z)$

$$z = \sum_k w_k x_k + b$$

b

Input layer   Hidden layers   Output layer

Figure 3. An artificial neuron and a neural network. (a) A single neuron takes inputs x_k and outputs $f\left(\sum_k w_k x_k + b\right)$, where f is called the activation function, and w_k's and b are weights and the bias of the neuron. (b) A neural network is composed of artificial neurons stacked in layers and connected to each other.

Figure 4. Results of an instance of the experiment for different initial states. Photon counts for all eight computation basis states over three qubits (000, 001, ..., 111) are binned into 30 μs time-bins in a total of 150 μs collection. The last column shows the total counts and darker color indicates more photons. The effect of crosstalk is visible in the intermediate channels (unlabeled rows).

and reducing the detection time. A readout fidelity of ∼99.99% has been demonstrated using this method [28].

Similar dark-to-bright or bright-to-dark pumping errors exist in other readout schemes, e.g., when a separate state with a finite lifetime is used as the dark state, known as a shelf state [7]. Smaller error contributions include laser light scattering off the ion trap and into the PMT as well as PMT dark counts, which account for 20 counts per second and 2 counts per second respectively. Both of these errors contribute only one false count for an average of 300 experiments and are therefore well discriminated using a thresholding method.

When detecting the state of more than one qubit, a bright ion can cause events on other ion detector channels. This crosstalk between the PMT channels modifies the distribution of observed photons, and the average detection fidelity decreases. One can choose a different threshold for each ion based on the state of its neighbors to partially mitigate these errors. In addition, using maximum likelihood methods, one can calculate the probability that an observed data point corresponds to the $|0\rangle$ or $|1\rangle$ state, and choose the most probable option [7, 29]. However, these methods are all tailored for a specific scenario and it is difficult to integrate other sources of information about the state, such as counts from extra PMT channels when imaging the ions onto alternating detectors, or photon arrival times. The latter contributes information about the state because bright-to-dark or dark-to-bright pumping events have characteristic photon arrival time distributions, i.e. photons arriving predominantly early or late in the detection window, which can be included in the discrimination procedure. To incorporate all data sources in a single framework and reduce the effect of crosstalk we take advantage of advances in the field of ML, and use an artificial neural network to perform the discrimination task. Before proceeding to the main results, we briefly introduce the neural network framework that is used in this work.

With N qubits, the measurement consists of photon counts and their arrival times on $M \geqslant N$ PMT channels. These photon counts can be binned into T time-bins to give $M \times T$ numbers that completely describe the measurements. Our goal is to classify these measurement results into 2^N states, in an N-qubit basis. Therefore, we consider a supervised learning scenario, where a set of measurement results and their corresponding states is used to train the machine and
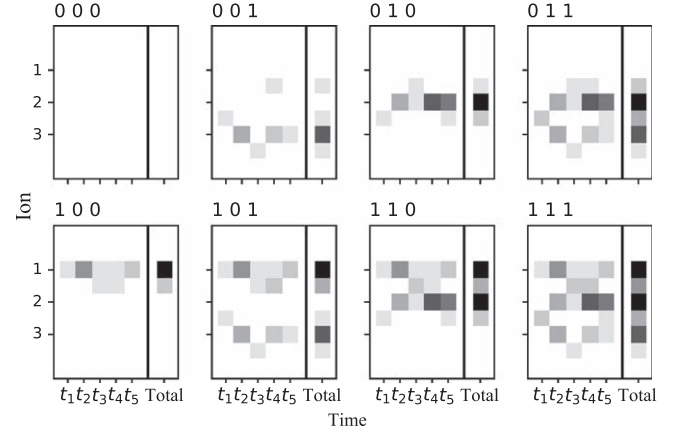
predict the correct state corresponding to a given input. We use a feed-forward neural network as depicted in figure 3 [30]. The network is built from a collection of neurons arranged in layers (columns in figure 3). A neuron is a unit that takes the input values x_k, and evaluates $f\left(\sum_k w_k x_k + b\right)$, where f is a mathematical function called the activation function, and w_k and b are scalar parameters that are referred to as neuron's weights and bias, respectively. There are various choices for the function f. A common example is the rectifier $f(z) = \max(0, z)$. When neurons are arranged in layers, the output of the previous layer serves as the input to the next layer. We represent weights and biases of the ith layer by the matrix $\mathbf{W}^{(i)}$ and the vector $\mathbf{b}^{(i)}$, respectively. In this notation, the rows of $\mathbf{W}^{(i)}$ represent the weights of the neurons in layer i, and elements of the vector $\mathbf{b}^{(i)}$ are the corresponding biases. Therefore, we have $\mathbf{x}^{(i+1)} = f(\mathbf{z}^{(i)})$, where $\mathbf{z}^{(i)} = \mathbf{W}^{(i)}\mathbf{x}^{(i)} + \mathbf{b}^{(i)}$, and the function f is applied element-wise. The first layer is called the input layer, where the neurons output the input data. Here, we have $M \times T$ neurons representing integrated photon counts from each ion in a time-bin (pixel values in figure 4). The last layer of the network is called the output layer. This layer captures different classes (states) that the input can take. Here, we use one-hot encoding to represent the classes n, corresponding to the 2^N different quantum states in our system (image labels in figure 4). That is, for each sample s, $y_{s,n} = 1$, if the sample is prepared in the nth state, and $y_{s,n} = 0$, otherwise. For classifying data into exclusive classes, it is common to use the softmax activation at the output layer, i.e. $f(\mathbf{z})_n = \frac{e^{z_n}}{\sum_k e^{z_k}}$. With this choice of activation, the output is normalized and can be interpreted as the probability of the input belonging to class n. The output with the highest probability is chosen as the predicted class. All the layers between the input and output are called hidden layers, and we use the rectifier function for them.

In order to predict the correct class, i.e., the quantum state, associated with each input, the network has to be trained. For

J. Phys. B: At. Mol. Opt. Phys. **51** (2018) 174006

A Seif *et al*

example, with 3 ions and 7 PMT channels, the machine should predict that the input $\mathbf{x}_s^{(in)} = (1\ 10\ 2\ 1\ 1\ 10\ 1)^T$ corresponds to the state $|101\rangle$, which is the 6th class and is represented by $\mathbf{y}_s = (0\ 0\ 0\ 0\ 0\ 1\ 0\ 0)^T$. The performance of the network can be quantified by a cost-function that measures the difference between the network's prediction and the target. Hence, the task of training is to find the weights and biases that optimize this cost-function. In this work, we use the ADADELTA optimizer [31] to minimize the cross entropy

$$C(\{\mathbf{W}^{(i)}\}_i, \{\mathbf{b}^{(i)}\}_i) = -\sum_{s=1}^{N_s}\sum_{n=1}^{2^N} y_{s,n}\log(\tilde{y}_{s,n})$$
$$+ (1 - y_{s,n})\log(1 - \tilde{y}_{s,n}), \qquad (1)$$

where N_s is the number of samples, 2^N is the number of classes, and $\tilde{y}_{s,n}$ is the output value of the network corresponding to the class n for the sample s.

To design and train the network, we split the data in three sets: training (60%), cross-validation (20%), and test sets (20%). The training set is used to train a given network and find its optimal weights and biases. The cross-validation set is used to evaluate the performance of networks with different number of hidden layers and neurons to choose the optimal network architecture. We observe that networks with two hidden layers perform the best. For each network, all hidden layers have the same number of neurons. This number varies from 8, for the simplest case, to 40, for the network with the most features. Such an architecture is complex enough to correctly classify the measurement data without overfitting it. Lastly, the reported performance of the optimal networks is evaluated, using the previously unseen test set.

We now discuss the results in detail. We begin by moving a single trapped-ion to the positions that would be occupied by ions in the multi-ion chain that we wish to investigate. This method allows us to to recreate the experimental setup with N qubits. We typically image ions onto alternating PMT channels to reduce the crosstalk, which leaves the intermediate channels unused. We also take data imaging them onto neighboring PMT channels in order to explore how detection errors would change for a chain of ions with smaller inter-ion distance. Then, we either initialize the ion in $|0\rangle$ to take data on dark states or we use a high-fidelity microwave pi-pulse to create $|1\rangle$ for bright-state data. Finally, we detect the qubit state by counting how many photons are detected on the ion's corresponding PMT channel as well as neighboring channels. In addition, the photon arrival time is recorded with sub-μs resolution. By loading only a single ion, we can create the full statistics for all the 2^N computational basis states by superimposing these individual distributions. This procedure separates SPAM errors from other systematic errors present in the system such as addressing crosstalk errors. Additional background noise from superimposing the statistic of individual qubits do not significantly contribute to errors. The average detection fidelity, which includes a small error from state preparation, is given by

$$\bar{\mathcal{F}} = 1/\mathcal{N}\sum_i p(\text{measured } i|\text{prepared } i), \qquad (2)$$

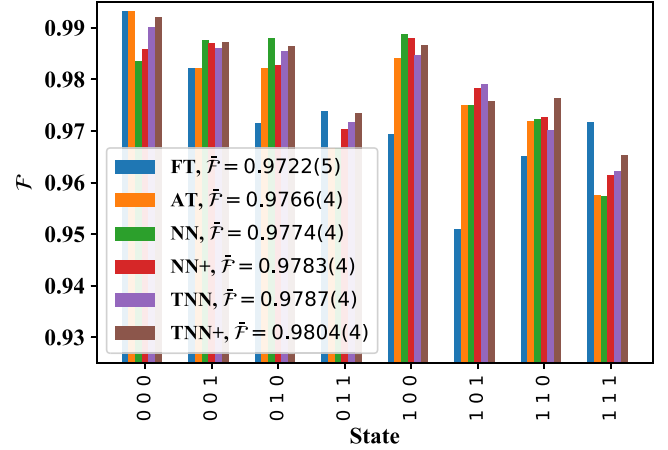where the sum is carried over all the computational basis states.

Figure 5. Comparison of different methods for state detection defined in the main body. We can see that the neural network (NN) methods outperform the conventional thresholding (FT, AT) methods. In addition, the performance is enhanced gradually as we provide the neural network with more features, e.g. intermediate channels and time stamps. The errors given in parentheses are statistical.

We compare six different methods and show that ML approaches outperform the two commonly used strategies in state discrimination. Below we describe these six strategies:

(i) Fixed threshold (FT): a threshold for photon counts is chosen to maximize the discrimination between bright and dark probability distributions. The same threshold is used for all the ions. In experiments with more than one qubit, this threshold is higher than the single qubit case because of crosstalk.

(ii) Adaptive threshold (AT): the threshold for each ion depends on the state of its neighbors. First, the state is determined by a fixed threshold, and then the inference process is iterated based on the state of neighbors and the corresponding thresholds.

(iii) Neural network (NN): first the photon counts from the ion PMT channels and their corresponding 2^N states (classes) are fed into a neural network. After the training, the neural network can predict the state of a given array of photon counts.

(iv) Neural network with intermediate channels (NN+): similar to NN, but the input also contains the intermediate PMT channel's data.

(v) Neural network with time-stamped data (TNN): the photon counts from the ion PMT channels are collected into time-bins to form a 2D image, where one axis is the time, and the other represents the location of the ions. The color intensity then represents the number of photons observed in that time-bin (see figure 4). These images with their corresponding labels are used to train the neural network.

(vi) Neural network with time-stamped data and intermediate channels (TNN+): the time-binned photon counts of the ion PMT channels and the intermediate channels are used to form an image, which the neural network learns
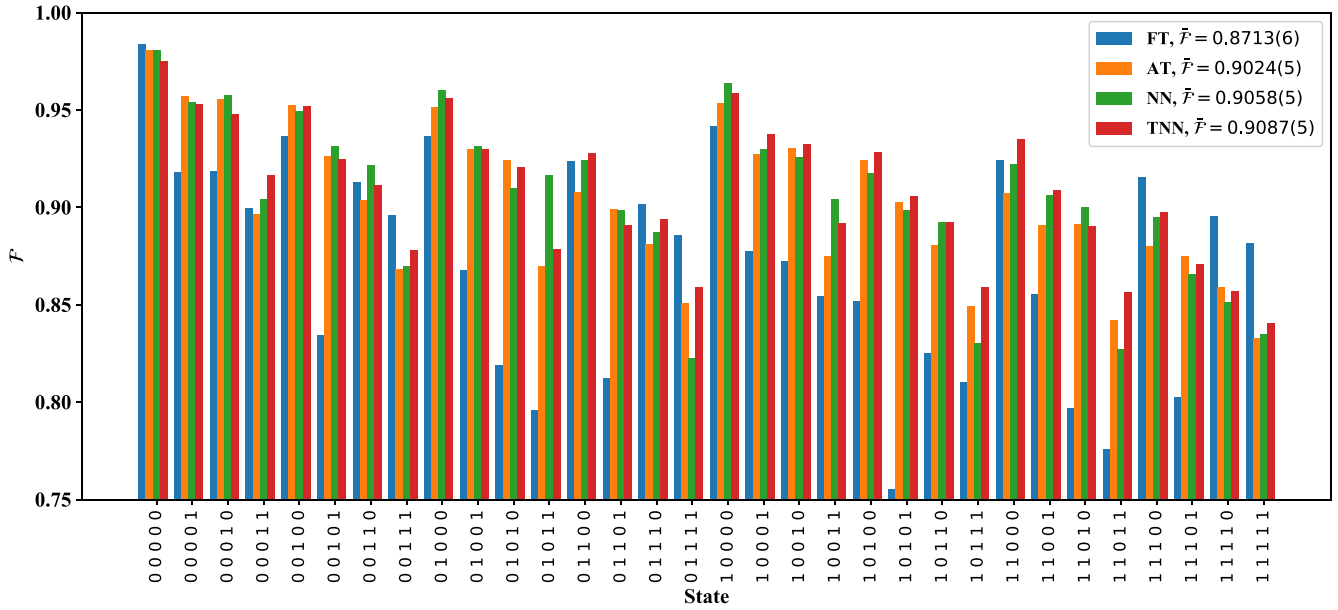
Figure 6. Comparing neural networks with threshold methods for five-qubit state detection. In this case the intermediate channel data is not available, but neural networks can still perform better than threshold methods. The errors given in parentheses are statistical.

to classify. This is the most comprehensive information available about the experiment.

We note that due to a large overlap between the photon count distributions of the intermediate channels with different bright neighbors, it is not possible to utilize the intermediate channel data with a simple threshold method. The same is true for the time-binned data, where the overlap of the bright and dark distributions is significant. This is because the distributions are Poissonian and have close mean values. However, the neural network can easily incorporate all the features and extract the available information.

In the first experiment, we consider a three qubit measurement scenario where the data from intermediate channels is available. We collect 80 000 samples for each class and apply the six strategies and observe that the neural network outperforms the other methods. In figure 5, it can be seen that with the same amount of information NN outperforms FT and AT, and when additional information is provided TNN+ can improve the errors over FT and AT by 30% and 17% respectively. It can also be seen that the neural network reduces the false positives in detecting 000 and 111 states and improves the crosstalk errors in the other states. Note the architecture of the neural network is kept the same and only the number of neurons are increased to represent the more complicated features, therefore providing a flexible tool for inferring properties of the system from experimental data.

In the second experiment, the ions are moved closer to each other to represent experiments where there are many ions in the trap, and neighboring PMT channels are associated with different ions. In this case, the data from intermediate channels is no longer available, and the crosstalk errors are increased. We consider a five-qubit measurement scenario with 50 000 samples for each class and compare FT, AT, NN and TNN methods.
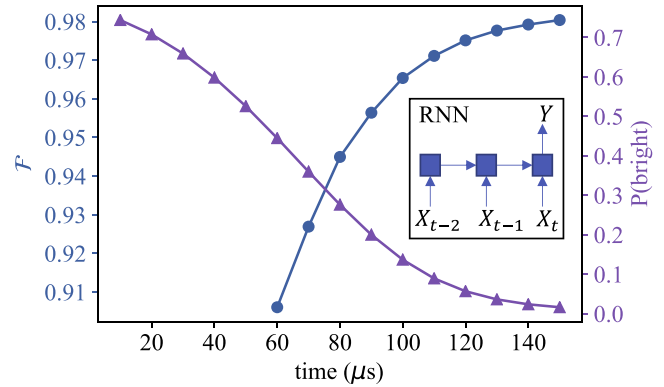

Figure 7. Recurrent neural network approach. The inset shows a schematic representation of the network, where the carried internal state and the output is affected by additional sequential inputs (time-binned photon counts). Left ordinate, blue circles: performance of a recurrent network for different detection times. The fidelity increases with longer detection times. Right ordinate, purple triangles: the average probability of the ion being bright decreases with the arrival time of the first photon.

As shown in figure 6 we can see that the same behavior observed in the first experiment persists, and neural networks beat threshold methods, and incorporating time-stamped data further improves the fidelities. Specifically, we observe 29% and 6% improvement by TNN over FT and AT, respectively.

In addition, we employ a recurrent neural network (RNN) as an alternative approach. These networks are tailored towards studying sequences of data (time-bins in our case), where the output in each step depends on the history through the internal state and an external input (see figure 7 inset). This feedback and memory effect is useful in capturing correlations in the sequence. While we observe the same fidelity as TNN+, this method is advantageous for experiments with variable detection time, since it can handle data sequences with different lengths.

We illustrate this capability by training a long short-term memory network, which is a type of RNN [32], with the full sequence of measurement data, using finer time-bins of 10 μs. Then, we evaluate the performance of the network by varying the length of the test sequence, and observe that indeed the performance increases with the measurement time (see circles in figure 7). In addition, we interrogate the network with artificial data to map out its internal mechanism. Specifically, we construct sequences with a single photon count, the arrival time of which is scanned. The output indicates the significance of the photon arrival time in deciding the state of the qubits. We observe that the network learns that photons with late arrival times are more likely to come from ions prepared in the dark state, which is consistent with our physical understanding of error mechanism by off-resonant excitation (see triangles in figure 7).

We have shown that a simple neural network classifier can improve the detection fidelities over tradition thresholding methods. The neural network classifier does not require assumptions about the system and can incorporate different data sources in one framework in a straightforward way. As the ion trap systems are very clean and the measurements are well-described theoretically, we do not expect neural networks to beat complicated models that take into account possible errors and evaluate the likelihood of a state corresponding to measured values. Similarly, we were not able to observe significant improvements over feed-forward networks by using more sophisticated techniques such as RNNs or convolutional neural networks. This is because the patterns and correlations in the data are simple and hence well-captured by feed-forward networks. However, we expect such advanced techniques to be especially useful in other systems like superconducting qubits, where the measurement processes are more complicated and the data has intricate features.

In addition, we have considered exclusive classes in our classifier, which implies the size of the network scales exponentially with the number of qubits. However, we have observed that multi-label classifiers can achieve a performance close to our method, while maintaining a linear scaling with the number of qubits. Moreover, while it is not necessary in the current setup, by taking advantage of the locality of the crosstalk one can bootstrap smaller networks over a few qubits. By taking majority vote over the outcome of the smaller classifiers, the most probable state corresponding to the measurement results over many qubits is decided.

In conclusion, we expect that techniques such as the one presented can simplify and improve future experiments and serve as a straightforward alternative for optimizing the readout of quantum computers as they are scaled up to many qubits.

Acknowledgments

ORCID iDs

Alireza Seif ⓘ https://orcid.org/0000-0001-5419-5999

References

[1] Shor P W 1995 *Phys. Rev.* A **52** R2493
[2] Steane A 1996 *Proc. R. Soc.* A **452** 2551
[3] Gottesman D 2009 arXiv:0904.2557
[4] Myerson A, Szwer D, Webster S, Allcock D, Curtis M, Imreh G, Sherman J, Stacey D, Steane A and Lucas D 2008 *Phys. Rev. Lett.* **100** 200502
[5] Jeffrey E *et al* 2014 *Phys. Rev. Lett.* **112** 190504
[6] Olmschenk S, Younge K, Moehring D, Matsukevich D, Maunz P and Monroe C 2007 *Phys. Rev.* A **76** 052314
[7] Burrell A H 2010 High fidelity readout of trapped ion qubits *Ph.D. Thesis* University of Oxford
[8] Jordan M I and Mitchell T M 2015 *Science* **349** 255
[9] Nielsen M A 2015 Neural networks and deep learning Determination Press http://neuralnetworksanddeeplearning.com
[10] Carleo G and Troyer M 2017 *Science* **355** 602
[11] Saito H 2017 *J. Phys. Soc. Japan* **86** 093001
[12] Gao X and Duan L-M 2017 *Nat. Commun.* **8** 662
[13] Deng D-L, Li X and Sarma S D 2017 *Phys. Rev.* X **7** 021021
[14] Deng D-L, Li X and Das Sarma S 2017 *Phys. Rev.* B **96** 195145
[15] Torlai G, Mazzola G, Carrasquilla J, Troyer M, Melko R and Carleo G 2018 *Nat. Phys.* **14**
[16] Torlai G and Melko R G 2018 *Phys. Rev. Lett.* **120** 240503
[17] Baireuther P, O'Brien T E, Tarasinski B and Beenakker C W 2018 *Quantum* **2** 48
[18] Krastanov S and Jiang L 2017 *Sci. Rep.* **7** 11003
[19] Torlai G and Melko R G 2017 *Phys. Rev. Lett.* **119** 030501
[20] Maskara N, Kubica A and Jochym-O'Connor T 2018 arXiv:1802.08680
[21] van Nieuwenburg E P, Liu Y-H and Huber S D 2017 *Nat. Phys.* **13** 435
[22] van Nieuwenburg E, Bairey E and Refael G 2017 arXiv:1712.00450
[23] Carrasquilla J and Melko R G 2017 *Nat. Phys.* **13** 431
[24] Wetzel S J 2017 *Phys. Rev.* E **96** 022140
[25] Magesan E, Gambetta J M, Córcoles A and Chow J M 2015 *Phys. Rev. Lett.* **114** 200501
[26] Debnath S, Linke N M, Figgatt C, Landsman K A, Wright K and Monroe C 2016 *Nature* **536** 63
[27] Acton M, Brickman K-A, Haljan P, Lee P, Deslauriers L and Monroe C 2006 *Quantum Inf. Comput.* **6** 465
[28] Crain S G 2016 Integrated system technologies for modular trapped ion quantum information processing *Ph.D. Thesis* Duke University
[29] Shen C and Duan L 2012 *New J. Phys.* **14** 053053
[30] LeCun Y, Bengio Y and Hinton G 2015 *Nature* **521** 436
[31] Zeiler M D 2012 arXiv:1212.5701
[32] Gers F A, Schraudolph N N and Schmidhuber J 2002 *J. Machine Learn. Res.* **3** 115