Mode-resolved thermometry of a trapped ion with deep learning

Yi Tao,* Ting Chen,* Hongyang Wang, Jie Zhang, Ting Zhang, Yi Xie,† Pingxing Chen, and Wei Wu‡

Institute for Quantum Science and Technology, College of Science, National University of Defense Technology, Changsha Hunan 410073, China;

Interdisciplinary Center for Quantum Information, National University of Defense Technology, Changsha Hunan 410073, China and Hunan Key Laboratory of Mechanism and Technology of Quantum Information, Changsha Hunan 410073, China

(Received 9 February 2024; revised 28 April 2024; accepted 23 May 2024; published 25 June 2024)

In a trapped-ion system, accurate thermometry of ions is crucial for precisely evaluating the system state and performing quantum operations. However, when the motional state of a single ion is far away from the ground state, the spatial dimensionality of the phonon state sharply increases. Then it is difficult to realize accurate and mode-resolved thermometry with existing methods. In this work, we apply deep learning to the thermometry of a trapped ion and propose an efficient and mode-resolved method for accurately estimating large mean phonon numbers. We have also conducted experimental verification based on a self-developed surface trap. The result has shown the significant accuracy and efficiency of the method for thermometry of a single ion with large mean phonon numbers. The mode-resolved feature of our method makes it better applied to the characterization of system parameters, such as evaluating cooling effectiveness and analyzing surface trap noise. Our trained neural network model can be easily extended to other experimental setups and parameter ranges without modification of the experimental setups.

DOI: 10.1103/PhysRevA.109.062434

I. INTRODUCTION

Trapped ions have emerged as one of the most promising platforms for quantum computing, primarily attributed to their exceptionally long coherence time [1] and the precise controllability of both internal and external states [2]. To thoroughly assess the performance of ion traps and facilitate coherent operations, it is crucial to accurately measure and control the effective temperature of ions.

For a single trapped ion near the motional ground state, there exist several well-established techniques for thermometry, including the sideband ratio method [3], singular value decomposition [4], etc. But these methods become impractical when the motional state is far from the ground state. While in scenarios such as heating-rate measurement [5,6] and dissipative-state preparation [7–9], thermometry is still necessary within a broader range of the mean phonon number $\bar{n}$. To address this, alternative approaches, e.g., dark resonances [5,10], spatial thermometry [11], and Doppler recooling [6] have been proposed. Dark resonances and spatial thermometry offer a measurement range from several hundred to 10 000 phonons, while Doppler recooling is specifically designed for $\bar{n} \geqslant 10^4$. Nevertheless, all of these techniques suffer from measurement complexity or inaccuracy [5,12].

More significantly, in many situations, it is necessary to separately measure the mean phonon number of individual vibration modes. For a single ion, different vibration modes mainly refer to distinct vibration frequencies from three motional directions. The aforementioned methods can only provide an estimation of the ensemble effective temperature determined by all relevant modes, instead of reliable information of a single vibration mode. In more detail, the dark resonances method offers fast probing of velocity distribution, which is based on the modifications of the photon scattering rate induced by the Doppler effect, resulting in vibration modes that are indistinguishable. In $^{40}Ca^{+}$ experiments, the 397- and 866-nm lasers need to be simultaneously controlled to have the same detuning to the upper-energy level, which increases experimental complexity and uncertainty. The spatial thermometry method makes use of the extent of ion position images. Therefore, the accuracy relies on the knowledge of the point spread function of the imaging system. This method also cannot provide information on each vibration frequency. Another method [13] using spin-dependent kick provides accurate thermometry for $\bar{n}$ in $[0.1, 10^4]$. The experiment involves a near-resonance Ramsey process. It requires a mode-locked laser and exact Rabi frequency, thereby limiting its scalability. Then the method still cannot distinguish separate vibration modes in the direction of the wave vector.

To tackle the challenges, we employ deep neural networks (DNNs) to estimate the mean phonon number of each single vibration mode of a single ion. Deep learning [14] excels in analyzing intricate structures of high-dimensional data and extracting meaningful features from it. In the field of natural sciences, deep learning has found diverse applications involving processing large amounts of complex data, such as identifying potential drug molecules [15] and analyzing biological data [16]. Particularly, its application in physics

*These authors contributed equally to this work.
†Contact author: xieyi2015@nudt.edu.cn
‡Contact author: weiwu@nudt.edu.cn

 ©2024 American Physical Society

[17,18] has mainly focused on particle physics [19], quantum many-body physics [20], and cosmology [21].

In this work, we present an accurate, efficient, and mode-resolved method based on DNNs for the thermometry of a single ion's motional state of $\bar{n}$ up to 1500, with no need for the precise value of coupling strength. The data to be measured only involve excited state populations of multiple orders of blue sidebands [22–24], determined by the Lamb-Dicke parameter η, the mean phonon number $\bar{n}$, the coupling strength Ω, and the pulse duration t. Because the carrier transition which involves joint coupling with all existing vibration modes is excluded, our method can avoid influence from other vibration modes, thus enabling mode-resolved thermometry.

The effective regime of our method spans from tens of phonons to over 1000 phonons. The range of the Lamb-Dicke parameter is between 0.07 and 0.21, allowing for the tunability of the measurable range of vibration frequency. We generated test datasets for theoretical testing, and the average error is less than 1% on the noise-free test set and still below 10% after introducing binomially distributed projection noise. Furthermore, we performed experimental verification of the method. The error of the best model using 15 blue sidebands was found to be less than 10%. Also, the experimental operation of our thermometry method is quite straightforward. The results can be promptly obtained through the utilization of deep neural networks, with no need for additional postprocessing.

II. METHOD

A. Theory and training set

For a single ion trapped in a harmonic potential, which remains a valid assumption for the larger $\bar{n}$ range concerned here [6,13], we only consider two internal states of the ion and the axial harmonic vibration with frequency ω_z. The state of the ion can be represented as $|\uparrow, m\rangle$ or $|\downarrow, n\rangle$, where m and n represent phonon numbers of the Fock state. The Rabi frequency $\Omega_{m,n}$ of the transition between these two states is given by [23,24]

$$\begin{aligned}\Omega_{m,n} &= \Omega|\langle m|e^{i\eta(a+a^\dagger)}|n\rangle| \\ &= \Omega e^{-\frac{\eta^2}{2}}\eta^{m-n}\sqrt{\frac{n!}{m!}}L_n^{m-n}(\eta^2),\end{aligned} \tag{1}$$

where we suppose $m \geqslant n$, and Ω is the coupling strength, L_n^α is the Laguerre polynomial, and η is the Lamb-Dicke parameter $\eta = \cos(\theta)k\sqrt{\hbar/2M\omega_z}$, which involves the ion mass M, the wave number k, the vibration frequency ω_z, and the angle θ between the wave vector and the vibration direction.

Assuming that the ion is initially prepared in the state $|\downarrow, n\rangle$, then weakly coupled to a blue-detuned laser, i.e., $\Omega \ll \omega_z$, the population of the $|\uparrow, m\rangle$ state would be $P_{|\uparrow,m\rangle} = \sin^2(\Omega_{m,n}t)e^{-\gamma_n t}$ after an evolution duration t. Ulteriorly, when the motional state of the ion is described by a density matrix with diagonal elements being the probability p_n of $|n\rangle$, the excited state population of the qth blue sideband will be

$$P_\uparrow(q) = \sum_n p_n \sin^2(\Omega_{n+q,n}t)e^{-\gamma_n t}. \tag{2}$$

The relaxation effect term $e^{-\gamma_n t}$ could be omitted, as the pulse duration t in our method is much shorter than the energy-level lifetime. Then $P_\uparrow(q)$ is rewritten as

$$P_\uparrow(q) = \sum_n p_n \sin^2(k_{n+q,n}\Omega t), \tag{3}$$

where $\Omega k_{n+q,n} = \Omega_{n+q,n}$. With the knowledge of η, p_n, and Ωt, the excited state populations of each blue sideband can be obtained using Eqs. (1) and (3) straightforwardly. However, when the mean phonon number is relatively large, it is challenging to determine the accurate value of Ω experimentally by fitting the Rabi oscillating curve. Due to the exceedingly small differences between $\Omega_{n+q,n}$ and $\Omega_{n'+q,n'}$ as $\bar{n}$ increases, it is difficult to get the exact value of $\bar{n}$. Furthermore, in Eq. (3), the upper limit of n needs to be high enough to obtain reliable accuracy. For a thermal state with $\bar{n} = 1500$, the upper limit of n should reach 6000 to cover 98% of the whole distribution, resulting in a quite time-consuming fitting process.

The essence of the above problem lies in the high dimensionality of the phonon state of large $\bar{n}$. Thanks to the capability of DNNs to process complex data, this situation lends itself to a suitable application scenario of the deep learning method. The initial step in constructing a DNN model involves generating a reliable training set. Typically, the dataset can be derived from public databases, experimental data, or theoretical calculations. However, in this application, due to the requirement of the precision and quantity of the training dataset, theoretical calculation is practically feasible. Here, we narrow our focus down to the thermal phonon state described by the probability distribution $p(m, \bar{n}) = \bar{n}^m/(\bar{n}+1)^{m+1}$. Considering our experimental configuration, we carefully choose discrete integer values ranging from 1 to 1500 as the variable $\bar{n}$. Furthermore, we restrict η in the range of 0.069 to 0.217, specifically tailored for the $^{40}Ca^+$ ion used in our system. This choice stems from the fact that ω_z falls within the range of 200 to 2000 kHz when the laser's wave vector aligns parallel to the direction of vibration. Additionally, we have empirically determined that the interval of 0.5π to 2π for Ωt can provide good pulse-time tolerance when the exact Ω is uncertain in experiments. Under various combinations of these parameters, we calculated $P_\uparrow(1)$ to $P_\uparrow(15)$. Each training data point comprises input values η and $P_\uparrow(1), P_\uparrow(2), \ldots, P_\uparrow(Q)$, and outputs $\bar{n}$ and Ωt. To ensure successful training, we have generated a dataset containing 1×10^6 evenly distributed samples, encompassing diverse parameter combinations.

In practical experiments, the electronic shelving technology allows for the identification of either a bright state or a dark state in each independent measurement. As a result, the measured probability $P_{k,N}$ follows a binomial distribution:

$$P_{k,N} = \frac{k}{N} = C_N^k P_\uparrow^k (1-P_\uparrow)^{N-k}, \tag{4}$$

where N represents the total number of measurements while k refers to the count of excited states. Smaller N brings greater noise, so the model should exhibit robustness against such noise for a limited N.

Figure 1 illustrates the architecture of the DNNs we employed. By training with data involving different numbers of blue sidebands, we have developed four distinct models, namely, M-5, M-8, M-10, and M-15, respectively involving sideband numbers $Q = 5$, 8, 10, and 15. Furthermore, to account for projection noise, we have trained an additional

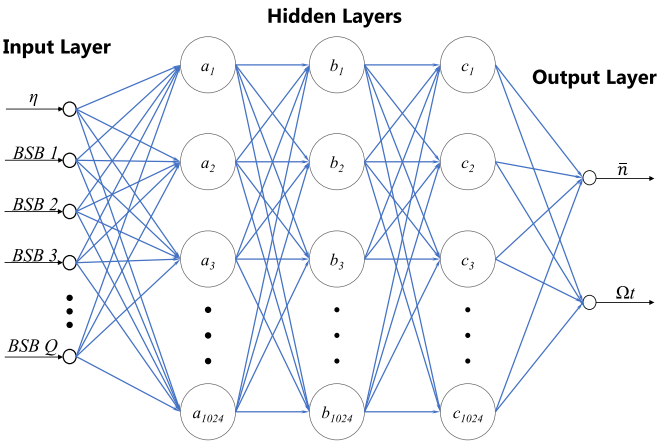

FIG. 1. DNN structure of our work. The input layer involves η and the first to the Qth order of the blue sideband; then there are three hidden layers each with 1024 neurons, and $\bar{n}$ and Ωt are two outputs.

model called ANM-5, which is based on M-5 and retrained by a noisy training set. In this noisy training set, each probability $P_\uparrow$ is transformed into $P_{k,N}$ through a simulated binomial distribution process. In this process, N is user defined and k is derived from $P_\uparrow$. It is worth noting that all of these models are trained by the same number of epochs and the same strategy.

B. Theoretical performance

To initially verify these models, we randomly generate 50 000 data items for testing, with $\bar{n}$ being a decimal number to ensure no overlap with the training set.

We first evaluated the performance of these DNN models with the test set under different mean phonon numbers. The results are illustrated in Fig. 2(a), where the relative error was $|\Delta\bar{n}|/\bar{n}$. For clarity, only 250 out of the original 50 000 points for each model are shown. When $\bar{n}$ is below 8, all models exhibit a relatively big error ($\geqslant$10%). This situation is primarily attributed to the fact that the excited state populations of some higher orders of the sideband approach 0, resulting in insufficient effective information. When $\bar{n} > 100$, the relative error of all models decreases to below 1%. When $\bar{n}$ approaches 1500, the relative error slightly increases. This increase is intrinsic to the limitations of deep learning, as data volume for training near the dataset boundary is relatively limited. Regarding the effects of utilizing different numbers of blue sidebands, M-15 and M-10 yield nearly identical results, while M-8 exhibits slightly inferior performance. Notably, there is an anomalous increase in error for M-5, which may stem from the reduced information on higher orders of sidebands.

Taking into account the inherent noise of the binomial distribution, we introduce different degrees of such noise into the test set by varying total numbers of measurements N. As illustrated in Fig. 2(b), each data point represents the average relative error calculated across the entire test set. Intuitively, the errors of all models decrease as noise decreases. Simultaneously, the error diminishes with an increasing number of sidebands, suggesting that employing more sidebands leads to enhanced robustness to the noise. Notably, ANM-5, which

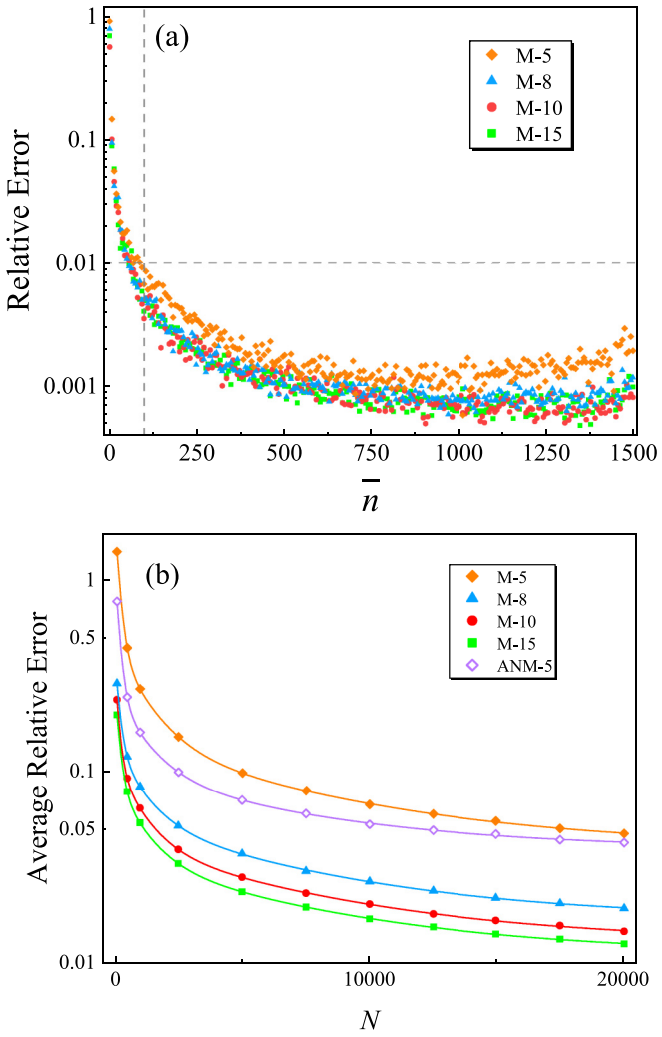

FIG. 2. Theoretical test results. (a) Relative error under different mean phonon numbers. There were 50 000 points for each model, here we only show 250 points to achieve a clearer presentation. When $\bar{n}$ is above 100, all models have an error of less than 1%, while an increased error is observed near the dataset boundary. (b) Average relative error on the entire test set with different total numbers of measurements N in Eq. (4). ANM-5 is trained based on M-5 using training data with binomial noise.

was trained using a noisy dataset, shows improved performance. The fitting curve reveals that, as N increases, the error initially decreases rapidly before reaching a point of diminishing returns. As N approaches infinity, the performance of ANM-5 does not significantly surpass that of the noise-free trained model, M-5, indicating that this training strategy still falls behind the gains achieved by increasing the number of sidebands.

The aforementioned analysis indicates that our method exhibits effectiveness for $\bar{n} \in [100, 1500]$. Its accuracy primarily depends on experimental measurement precision, since the noise-free test shows better results. To achieve an error of less than 5%, it would require thousands of measurements to obtain a single value of population. To avoid this in our experiment, peak fitting was utilized. The details are elaborated in the subsequent section.

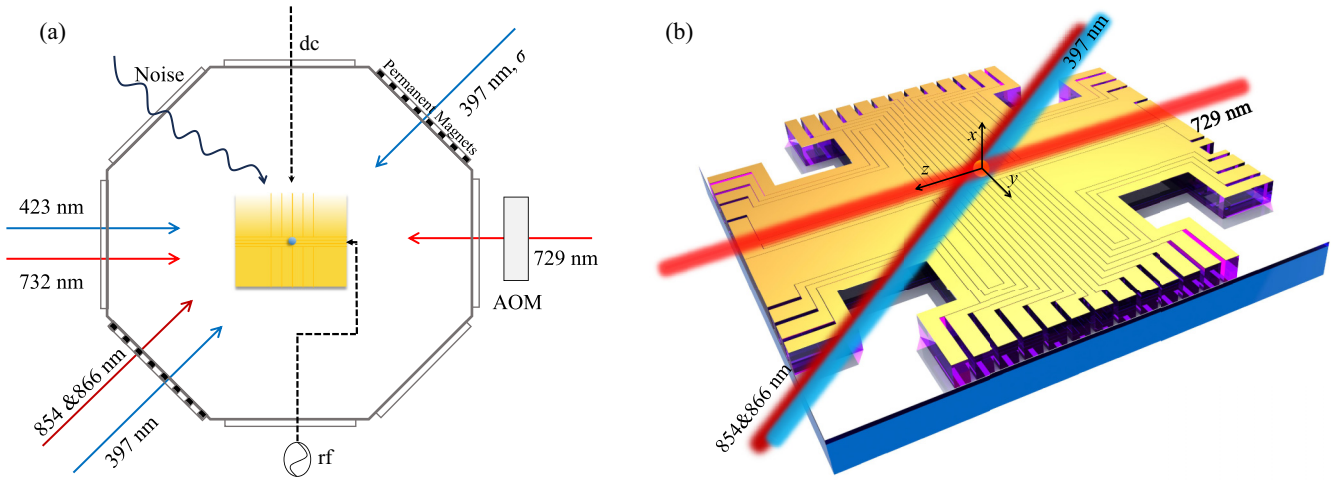

FIG. 3. Experimental setup. (a) Optical and electrical settings. The dashed lines represent dc and rf voltages loaded on electrodes of the trap, and the black curve represents artificially added noise. (b) Detailed model of the surface trap.

III. EXPERIMENTAL VERIFICATION

A. Experimental setup and scheme

The experiment was conducted on a six-wire surface trap as shown in Fig. 3, where the trap frequency could be adjusted flexibly and the target thermal state could be prepared efficiently. This trap has 15 pairs of dc electrodes, of which 7 pairs are mainly used to trap ions (Fig. 3(b)). A single $^{40}Ca^{+}$ ion is trapped at a distance of approximately 150 μm above the chip surface with the axial frequency ω_z set to 626 kHz, by applying a radio frequency voltage of 20.9 MHz and dc voltages. Other dc electrodes are given appropriate voltages to compensate for micromotion.

Two Zeeman sublevels of the $^{40}Ca^{+}$ ion, $^{2}S_{1/2}(m_j = -1/2)$ and $^{2}D_{5/2}(m_j = -5/2)$, are used as ground $|\downarrow\rangle$ and excited states $|\uparrow\rangle$. They are generated by a magnetic field whose direction is at an angle of 45° to the axial direction. The state is manipulated by a 729-nm laser parallel to the axial direction, whose switch and frequency are controlled by an acousto-optic modulator (AOM). The Lamb-Dicke parameter η is approximately 0.122, which is adjustable to change ω_z by changing the dc voltages. The cooling methods include Doppler cooling and sideband cooling. The 397- and 866-nm lasers are used for Doppler cooling, while the 729- and 854-nm lasers are subsequently used for sideband cooling to cool the ion to about 0.3 phonons. Under the above conditions, the measured heating rate is 0.245 ± 0.046 ms^{-1}. Due to the relatively low heating rate, and for experimental convenience, we additionally introduced a 5 V_{pp} noise signal to the 15th pair of electrodes to expedite the heating process.

To ascertain the effectiveness of our approach, it is crucial to accurately generate thermal states of various $\bar{n}$. With a good initial state and a constant heat rate [13], we can accurately prepare thermal states of any $\bar{n}$ we need. The experimental verification procedure is as follows.

(i) We measured $\bar{n}$ using the Rabi oscillation fitting method under various heating durations near the motional ground state. The fitting method was deemed accurate within the specified range, allowing us to obtain the accurate heating rate, and this artificial heating rate remained constant over time [3,5,13].

(ii) Based on the determined heating rate, subject the ion to an extended heating duration to attain the target thermal state with a larger mean phonon number. This provides an accurate value of $\bar{n}$ for reference.

(iii) Right after the heating, the frequency of the 729-nm laser was controlled to measure $P_{\uparrow}$ up to the 15th blue sideband. It is worth noting that, to improve efficiency and minimize measurement error, we do not scan such a large frequency range in Fig. 4(a). We actually conducted a small-frequency-range scan around each resonance point as shown

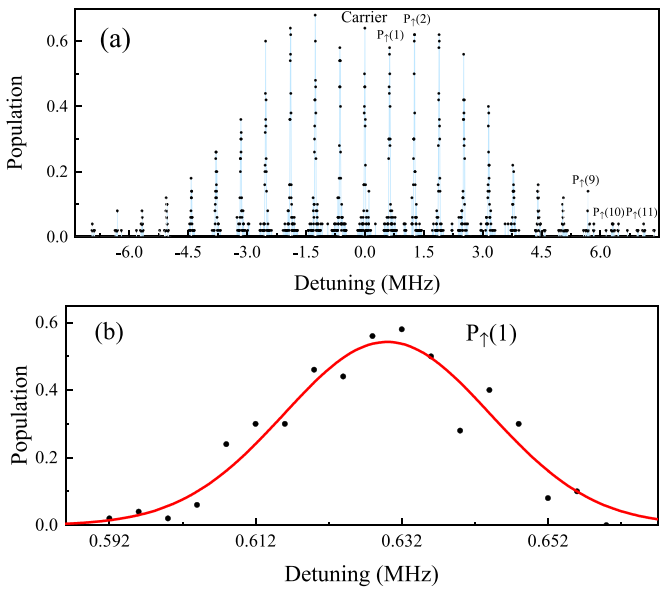

FIG. 4. (a) Excitation spectrum of motional sidebands. For higher orders of the blue sideband whose $P_{\uparrow} \leqslant 0.01$, it is viable to set these values to 0. (b) Gaussian fitting curve. In the practical experiment, there is no need to scan such a large range of frequency in panel (a), but rather finely scan small areas around each sideband and find the peaks with Gaussian fitting, by which frequency error and ac-Stark shift can be ruled out.

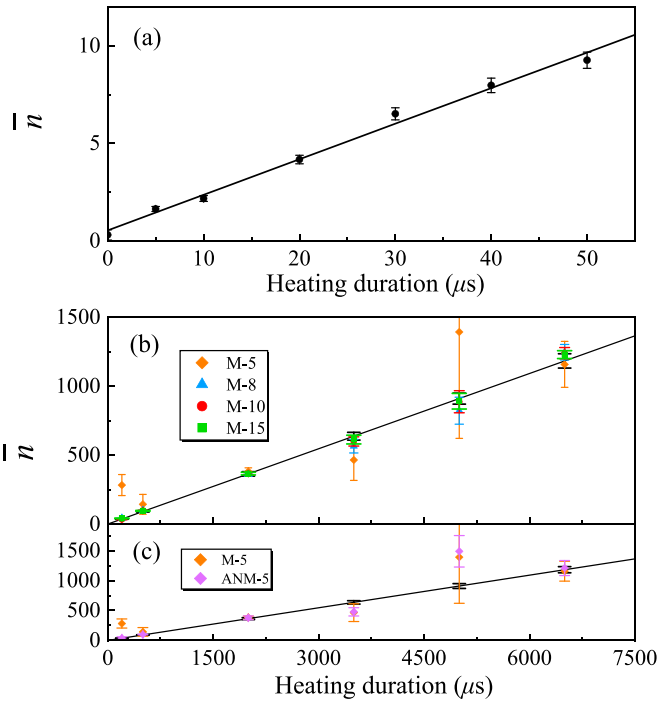

FIG. 5. Experiment results. (a) Heating rate measured by fitting the Rabi oscillation. (b) DNN result in a larger phonon number scale. (c) Comparison of M-5 and ANM-5. The black lines in panels (b) and (c) are the extrapolated heating line in panel (a).

in Fig. 4(b) and determined the value of the population by peak fitting using Gaussian distribution.

(iv) Upon completion of the experiment, we input the measured population values into the DNN model and compare the output with the reference value.

Due to the artificially introduced constant heating rate, this experimental scheme ensures the precision for generating thermal states with different mean phonon numbers and provides a reliable verification of our method.

B. Result

The heating rate measured by fitting the Rabi oscillation is shown in Fig. 5(a). By adding a noise signal, the heating rate is increased from 0.245 ± 0.046 to 182 ± 8 ms^{-1}, and we selected several heating duration points for verification.

Under different heating durations, thermal states of different $\bar{n}$ were obtained, and then the populations of multiple orders of sidebands for different thermal states were measured as input into DNN models. The output results are illustrated in Fig. 5(b). Within the range of mean phonon numbers spanning from tens of phonons to over 1000 phonons, models employing eight or more sidebands have comparable performance and demonstrate a commendable accordance with the extrapolated heating line. In contrast, M-5 exhibits relatively poor performance, largely attributed to its being unresistant to the severe noise encountered in real-world experiments. Yet, M-5 might yield improved results with more measurement times, given that its theoretical performance is acceptable in the above theoretical noise-free test.

The error bars in the figure were generated using Monte Carlo simulation [25], where we employed the uncertainties introduced during peak fitting as standard deviations in a normal distribution. Then a series of input values generated by simulating such normal distribution were fed into the DNN model to obtain a set of outputs whose standard deviation was taken as the error bar. Therefore, the size of error bars is influenced by both the accuracy of measurements and the inherent characteristics of the neural networks. Remarkably, an increased number of employed sidebands leads to a reduction in the size of the error bars, signifying that a greater number of sidebands contributes to enhanced robustness, which coincides with the theoretical test presented in the previous section. The above analysis suggests that using more sidebands brings better performance. However, since the difference among M-8, M-10, and M-15 is not significant, it is feasible to use M-10 for convenience in practical experiments.

Furthermore, ANM-5 was tested and shown in Fig. 5(c). Overall, training under projection noise resulted in a reduction in the error bar. In the low-$\bar{n}$ range, its performance surpassed that of M-5, although there was minimal difference in the higher range. The lack of substantial improvement, especially at the 5000-μs point, may be attributed to inherent defects in M-5 that persisted even after retraining with noisy data.

IV. SUMMARY

In this work, we have shown the application of deep neural networks to the thermometry of a single ion. Our method has demonstrated remarkable performance across a wide range of mean phonon numbers from tens of phonons to over 1000 phonons. Additionally, we explored the impact of employing different numbers of sidebands and revealed that an increased number of sidebands enhances both accuracy and robustness. The model trained with noisy data exhibits certain improvements in robustness but falls short of the gains achieved by increasing the number of sidebands. For practical applications balancing between efficiency and accuracy, ten sidebands are deemed sufficient within the above-mentioned $\bar{n}$ range.

Although this work only focused on the thermal distribution of a single ion's phonon state, the method's applicability can extend to various distributions in principle. When confronted with different ranges and distributions, creating tailored training datasets will suffice. It is worth emphasizing that the most advantageous feature of our DNN method is its mode-resolving capability. Furthermore, it can be generalized to mode-resolved thermometry of multiple ions [26]. Additionally, this method facilitates swift deployment across different experimental setups and ion types, requiring no modifications to the optoelectronic settings. Benefiting from the efficiency inherent in neural networks, results are available almost immediately right after the acquisition of experimental data.

In summary, we present an accurate, efficient, and mode-resolved approach for estimating the mean phonon number of a single trapped ion, particularly in the case of large $\bar{n}$. In the future, we envision extending this method to accommodate various scenarios such as the estimation of the phonon

coherent state, feedback control of system setups, and more crucially, the thermometry of multiple ions.

ACKNOWLEDGMENTS

This work is supported by the National Natural Science Foundation of China under Grants No. 12174447, No. 11904402, No. 12004430, No. 12074433, No. 12174448, and No. 12204543 and the Science and Technology Innovation Program of Hunan Province under Grant No. 2022RC1194.

APPENDIX: DESIGN OF DNNs

A typical DNN model usually consists of an input layer, some hidden layers, and an output layer. The dimensionalities of the input and output layers are determined by actual requirements, and parameters such as the number of hidden layers and the number of neurons in each layer are manually tuned. The structure of the network, training strategies, and the size and quality of the training set, all affect the effectiveness of the model. To build the DNNs, here we use PYTORCH [27], an open-source deep learning framework known for its excellent flexibility and ease of use.

The neural network in this work is composed of an input layer, three hidden layers, and an output layer. The population values as input naturally range from 0 to 1, making it suitable as input for the network with no need for postprocessing. There are 1024 neurons in each hidden layer. The activation functions are Tanh and ReLU [28].

For the large order of magnitude difference between the two output values, which is not conducive to evaluating the convergence effect and updating weights, the ranges of their values are scaled to the same interval: $[-10, 10]$. To evaluate the distance between the predicted value and the target value, we use the following loss function:

$$\text{loss}(x, y) = \frac{|y_1 - f(x)| + |y_2 - g(x)|}{2}, \tag{A1}$$

where x represents input, $f(x)$ and $g(x)$ are outputs, and

$$y_1 = \frac{\bar{n}}{75} - 10,$$
$$y_2 = \frac{15}{2}\left(\frac{\Omega t}{\pi} - \frac{1}{2}\right) - 10. \tag{A2}$$

Through multiple rounds of iteration, the parameters networks can be updated automatically by the optimizer Adam [29], until the loss converges. For example, the loss of M-15 converged to 0.009 after 20 epochs of training.

[1] C. D. Bruzewicz, J. Chiaverini, R. McConnell, and J. M. Sage, Trapped-ion quantum computing: Progress and challenges, Appl. Phys. Rev. **6**, 021314 (2019).
[2] C. J. Ballance, T. P. Harty, N. M. Linke, M. A. Sepiol, and D. M. Lucas, High-fidelity quantum logic gates using trapped-ion hyperfine qubits, Phys. Rev. Lett. **117**, 060504 (2016).
[3] F. Diedrich, J. C. Bergquist, W. M. Itano, and D. J. Wineland, Laser cooling to the zero-point energy of motion, Phys. Rev. Lett. **62**, 403 (1989).
[4] D. M. Meekhof, C. Monroe, B. E. King, W. M. Itano, and D. J. Wineland, Generation of nonclassical motional states of a trapped atom, Phys. Rev. Lett. **76**, 1796 (1996).
[5] V. Tugayé, J.-P. Likforman, S. Guibal, and L. Guidoni, Absolute single-ion thermometry, Phys. Rev. A **99**, 023412 (2019).
[6] R. J. Epstein, S. Seidelin, D. Leibfried, J. H. Wesenberg, J. J. Bollinger, J. M. Amini, R. B. Blakestad, J. Britton, J. P. Home, W. M. Itano, J. D. Jost, E. Knill, C. Langer, R. Ozeri, N. Shiga, and D. J. Wineland, Simplified motional heating rate measurements of trapped ions, Phys. Rev. A **76**, 033411 (2007).
[7] J. F. Poyatos, J. I. Cirac, and P. Zoller, Quantum reservoir engineering with laser cooled trapped ions, Phys. Rev. Lett. **77**, 4728 (1996).
[8] D. Kienzler, H.-Y. Lo, B. Keitch, L. De Clercq, F. Leupold, F. Lindenfelser, M. Marinelli, V. Negnevitsky, and J. P. Home, Quantum harmonic oscillator state synthesis by reservoir engineering, Science **347**, 53 (2015).
[9] Y. Lin, J. P. Gaebler, F. Reiter, T. R. Tan, R. Bowler, A. S. Sørensen, D. Leibfried, and D. J. Wineland, Dissipative production of a maximally entangled steady state of two quantum bits, Nature (London) **504**, 415 (2013).
[10] J. Roßnagel, K. N. Tolazzi, F. Schmidt-Kaler, and K. Singer, Fast thermometry for trapped ions using dark resonances, New J. Phys. **17**, 045004 (2015).
[11] S. Knünz, M. Herrmann, V. Batteiger, G. Saathoff, T. W. Hänsch, and Th. Udem, Sub-millikelvin spatial thermometry of a single Doppler-cooled ion in a Paul trap, Phys. Rev. A **85**, 023427 (2012).
[12] M. Brownnutt, M. Kumph, P. Rabl, and R. Blatt, Ion-trap measurements of electric-field noise near surfaces, Rev. Mod. Phys. **87**, 1419 (2015).
[13] K. G. Johnson, B. Neyenhuis, J. Mizrahi, J. D. Wong-Campos, and C. Monroe, Sensing atomic motion from the zero point to room temperature with ultrafast atom interferometry, Phys. Rev. Lett. **115**, 213001 (2015).
[14] Y. LeCun, Y. Bengio, and G. Hinton, Deep learning, Nature (London) **521**, 436 (2015).
[15] W. Patrick Walters and R. Barzilay, Applications of deep learning in molecule generation and molecular property prediction, Acc. Chem. Res. **54**, 263 (2021).
[16] V. I. Jurtz, A. R. Johansen, M. Nielsen, J. J. A. Armenteros, H. Nielsen, C. K. Sønderby, O. Winther, and S. K. Sønderby, An introduction to deep learning on biological sequence data: examples and solutions, Bioinformatics **33**, 3685 (2017).
[17] P. Sadowski and P. Baldi, Deep learning in the natural sciences: Applications to physics, in *Braverman Readings in Machine Learning: Key Ideas from Inception to Current State* edited by L. Rozonoer, B. Mirkin, and I. Muchnik (Springer, Berlin, 2018), pp. 269–297.

[18] G. Carleo, I. Cirac, K. Cranmer, L. Daudet, M. Schuld, N. Tishby, L. Vogt-Maranto, and L. Zdeborová, Machine learning and the physical sciences, Rev. Mod. Phys. **91**, 045002 (2019).

[19] D. Guest, K. Cranmer, and D. Whiteson, Deep learning and its application to LHC physics, Annu. Rev. Nucl. Part. Sci. **68**, 161 (2018).

[20] Alexandra Nagy and V. Savona, Variational quantum Monte Carlo method with a neural-network ansatz for open quantum systems, Phys. Rev. Lett. **122**, 250501 (2019).

[21] A. E. Firth, O. Lahav, R. S. Somerville, and R. S. Somerville, Estimating photometric redshifts with artificial neural networks, Mon. Not. R. Astron. Soc. **339**, 1195 (2003).

[22] D. J. Wineland, C. Monroe, W. M. Itano, D. Leibfried, B. E. King, and D. M. Meekhof, Experimental issues in coherent quantum-state manipulation of trapped atomic ions, J. Res. Natl. Inst. Stand. Technol. **103**, 259 (1998).

[23] K. E. Cahill and R. J. Glauber, Ordered expansions in Boson amplitude operators, Phys. Rev. **177**, 1857 (1969).

[24] D. J. Wineland and Wayne M. Itano, Laser cooling of atoms, Phys. Rev. A **20**, 1521 (1979).

[25] S. Raychaudhuri, Introduction to Monte Carlo simulation, in *2008 Winter Simulation Conference*, edited by S. J. Mason, R. R. Hill, L. Mönch, O. Rose, T. Jefferson, and J. W. Fowler (IEEE Xplore, Piscataway, NJ, 2008), pp. 91–100.

[26] Z. Jia, Y. Wang, B. Zhang, J. Whitlow, C. Fang, J. Kim, and K. R. Brown, Determination of multimode motional quantum states in a trapped ion system, Phys. Rev. Lett. **129**, 103602 (2022).

[27] A. Paszke, S. Gross, F. Massa, A. Lerer, J. Bradbury, G. Chanan, T. Killeen, Z. Lin, N. Gimelshein, L. Antiga *et al.*, PyTorch: An imperative style, high-performance deep learning library, in *Advances in Neural Information Processing Systems*, edited by H. Wallach, H. Larochelle, A. Beygelzimer, F. d'Alché-Buc, E. Fox, and R. Garnett (Curran Associates, Inc., New York, 2019), Vol. 32.

[28] X. Glorot, A. Bordes, and Y. Bengio, Deep sparse rectifier neural networks, in *Proceedings of the 14th International Conference on Artificial Intelligence and Statistics*, edited by G. Gordon, and D. Dunson, and M. Dudík (PMLR, New York, 2011), Vol. 15, pp. 315–323.

[29] D. P. Kingma and L. J. Ba, *Adam: A Method for Stochastic Optimization*, ICLR 2015 (Ithaca, New York, 2015).