Editors' Suggestion | Featured in Physics

Benchmarking a Tunable Quantum Neural Network on Trapped-Ion and Superconducting Hardware

Djamil Lakhdar-Hamina,[1,*] Xingxin Liu,[1] Richard Barney,[1] Sarah H. Miller,[2] Alaina M. Green,[1,3] Norbert M. Linke,[1,3,4] and Victor Galitski[1]

[1]Joint Quantum Institute and Department of Physics, University of Maryland, College Park, Maryland 20742, USA
[2]Applied Research Laboratory for Intelligence and Security, University of Maryland, College Park, Maryland 20742, USA
[3]National Quantum Laboratory (QLab), University of Maryland, College Park, Maryland 20742, USA
[4]Duke Quantum Center and Department of Physics, Duke University, Durham, North Carolina 27701, USA

(Received 27 August 2025; revised 6 February 2026; accepted 21 May 2026; published 22 July 2026)

We implement a quantum neural network on trapped-ion and IBM superconducting quantum computers for Modified National Institute of Standards image classification. Feedforward is realized through qubit rotations conditioned on measurement outcomes from previous layers. The network is trained classically, while inference is performed experimentally on quantum hardware. A tunable interpolation parameter connects the classical and quantum regimes. Moderate values improve classification performance by introducing measurement uncertainty. For borderline images misclassified classically but correctly identified quantum mechanically, we observe strong deviations from idealized simulations due to physical noise, consistent with fluctuations between nearby minima in the classification landscape. We further benchmark noise by inserting additional single- and two-qubit gate pairs into the circuits. Our results motivate more complex quantum neural networks on noisy intermediate-scale quantum devices and suggest a possible route toward classically nonsimulable architectures.

DOI: 10.1103/9bp2-42v3

Neural networks exhibit deep connections to spin models in statistical mechanics [1–4]. Quantizing neurons and weights into qubits and gates naturally gives rise to quantum neural networks [5]. Much recent work in quantum machine learning has focused on identifying quantum advantage [6–14]. However, few proposed architectures have been experimentally implemented on quantum hardware [15–17], particularly for benchmark image-classification tasks [18,19]. Moreover, most do not allow smooth interpolation between classical and quantum regimes, complicating direct comparisons between neural networks and quantum neural networks.

In the following, we present a quantum neural network architecture, the benchmark quantum neural network (BQNN), with several attractive features. First, it smoothly interpolates between classical and quantum regimes, reproducing a binary neural network in the classical limit. Second, it enables direct benchmarking on standard datasets and classical-quantum comparisons. Third, the circuit is deployable on current quantum hardware with modest qubit, fidelity, and connectivity requirements. Fourth, it generalizes to more complex architectures with partial midcircuit measurements and feedback, bridging quantum neural networks and measurement-induced phase transitions [20–22].

After reviewing the theory of this BQNN architecture, we present image-classification results on both superconducting and trapped-ion hardware. The experimental results on quantum devices are largely consistent with the expected behavior informed by classical simulations. In an intermediate quantum regime where quantum uncertainty from measurements is present but does not fully randomize inference the quantum modality outperforms the classical counterpart. We further study the effect of device noise and find that moderate physical noise does not hinder inference and can even enhance performance. Strong gate noise, modeled by inserting single- and two-qubit gate pairs and their inverses, eventually reduces performance to random guessing. To our knowledge, this is the first comparison of multiple quantum neural network hardware platforms on a benchmark image-classification task.

Reference [5] presents theory and simulation of a quantum partially binarized multilayer perceptron, i.e., the BQNN. This network may be smoothly tuned between classical and quantum modes of operation. This network is partially binarized in the sense that, while each training weight can take on any real value, activations of the hidden layers are constrained to be ± 1. For a classical network this is done by using the activation function

$$\phi(x) = \mathrm{sgn}(x). \tag{1}$$

The activation of the ith neuron in the kth layer is

*Contact author: dlakhdar@umd.edu

 © 2026 American Physical Society

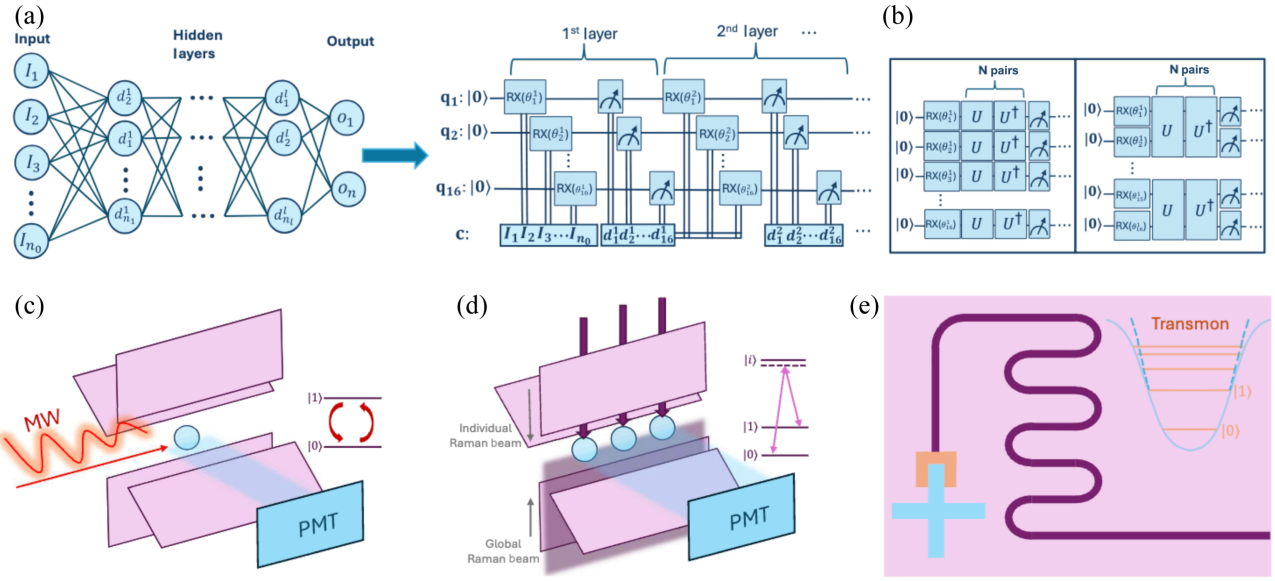


FIG. 1. (a) The mapping from the classical binarized multilayer perceptron network to the quantum circuit implementing the hidden layers of the BQNN. Each step of the forward pass consists of rotating each qubit by an angle controlled by the classical channels, then measuring each qubit. These measurements are the activations that are then passed to the next layer of the network. The networks consist of an input layer, l hidden layers, and an output layer. (b) The BQNN circuit with noise injection by means of inserted U and $U^\dagger$ pairs. (c)–(e) The varieties of quantum hardware on which we run the BQNN: trapped ions with microwave driven gates, trapped ions with laser driven gates, and superconducting qubits.

$$d_i^k = \phi\big((W^k \mathbf{d}^{k-1})_i\big), \tag{2}$$

where $\mathbf{d}^0 = \mathbf{I}$ is the input vector and each W^k is a matrix of trainable parameters. In this way, information travels through the network layer by layer in a process called feedforward. The output of the network with $\ell - 1$ hidden layers is $\mathbf{o} = W^\ell \mathbf{d}^{\ell-1}$, which is then the argument of the loss function during training.

This classical network is then quantized in a straightforward way: binary neurons are promoted to qubits and activation functions are implemented by single-qubit rotations followed by projective measurements. By defining the activations as the results of projective measurements, the activations remain binarized. The circuit which implements the hidden layers is shown on right hand side in Fig. 1(a).

To tune between quantum and classical regimes, we define the rotation angle of qubit i in layer k as

$$\theta_i^k = \frac{\pi}{2}\begin{cases} 1 - \phi_a\big((W^1 \mathbf{I})_i\big), & k = 1 \\ d_i^{k-1} - \phi_a\big((W^k \mathbf{d}^{(k-1)})_i\big), & k > 1 \end{cases}, \tag{3}$$

$$\phi_a(x) = \mathrm{htanh}(x/a) = \begin{cases} \mathrm{sgn}(x), & |x| \geq a \\ x/a, & |x| < a \end{cases}, \tag{4}$$

where ϕ_a may be considered a quantum activation function. In the limit $a \to 0$, ϕ_a reduces to the classical activation function Eq. (1). In this case, each rotation angle is either 0 or $\pm\pi$. This means that, under ideal conditions, the rotation gates will only move the qubits between computational basis states, and the behavior of the $a = 0$ network is deterministic and is equivalent to the standard classical neural net. Nonzero values of a enable arbitrary angels of rotation and hence nondeterministic measurements. We identify a as a parameter describing the "quantumness" of the network.

BQNN training was done using stochastic gradient descent with momentum [23] and a straight-through estimator to circumvent nondifferentiability of activation function [24].

Unlike deterministic classical inference, quantum classifications fluctuate between runs, requiring repeated sampling and majority-vote classification. Reference [5] explored both training and inference with a classically simulated BQNN on the Modified National Institute of Standards (MNIST) handwritten digit recognition dataset [25], finding improved classification performance in an intermediate quantum regime for $0 < a < \mathcal{O}(1)$ using around ~10 samples per image.

In the present Letter we train the network using classical simulation, but perform *classification* on two different sets of hardware, i.e., IBM circuit-QED processors based on transmon qubits [26,27], the trapped-ion quantum computer at the Joint Quantum Institute, with microwave gates and Raman gates. The network consists of three layers with 16 qubits per layer and is classically trained using a straight-through estimator on the MNIST dataset. While the circuit in Fig. 1(a) is constructed on 16 qubits according to the network's architecture, the circuit is separable and hence can be executed on as few as one qubit. Training and classification code can be found in the repositories [28,29].

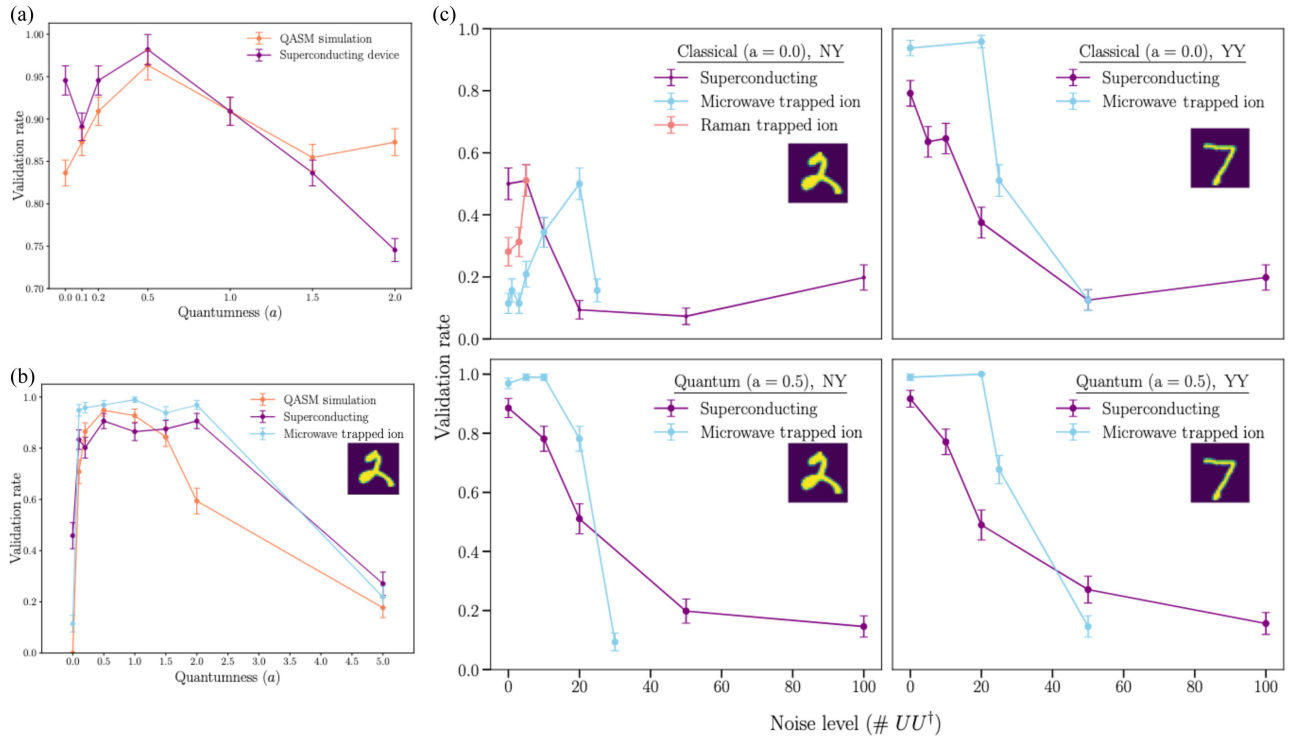


FIG. 2. (a) Validation rates on 55 randomly selected test images obtained from superconducting hardware and simulation for different values of a. For each image, ten shots were taken and the most common outcome was chosen as the predicted class. (b) An NY image that was misclassified at $a = 0.0$ shows improved classification accuracy when the quantumness a is increased to some optimal level on the three different platforms, i.e., superconducting, trapped-ion, and classical simulation. (c) For an NY and YY image, taken at either $a = 0.0$ (classical) and $a = 0.5$ (quantum), noise is controlled by the insertion of a number of single-qubit $UU^\dagger = I$ pairs.

The ion trap system is described in detail in Ref. [30]. It consists of $^{171}Yb^+$ ions confined in a linear Paul trap. Each ion possesses two types of quantum states: internal versus motional states. The internal states encode qubits and correspond to hyperfine levels within the ground-state manifold, $^2S_{1/2}$. Specifically, the qubit basis states $|0\rangle$ and $|1\rangle$ are defined as $|F, m_F\rangle = |0, 0\rangle$ and $|1, 0\rangle$, respectively. The motional states arise from quantized collective oscillations of the ion chain within the harmonic potential formed by the radio-frequency and static electric field. These vibrational modes serve as a quantum bus and mediate interactions between qubits.

We employ two trapped-ion gate architectures: microwave-driven single-qubit gates and laser-driven Raman transitions that, respectively, implement both single-qubit and two-qubit entangling gates. For Raman-based operations, three ions are confined simultaneously, though only left and right ions were used. Single-qubit gates are implemented using stimulated Raman transitions, driven by two beams: one globally illuminates the ion chain, while the other is split into focused beams for individual ion addressing. The frequency difference matches qubit splitting and drives coherent transitions. Entangling gates are implemented by applying two bichromatic laser beams per ion tuned near the red and blue sidebands of a shared motional mode. These beams generate spin-dependent forces that couple the ions' internal (spin) states through their collective motion, realizing high-fidelity Molmer-Sorensen interaction. For microwave gates, a single ion is trapped, and transitions between the states are driven using resonant microwave radiation. This approach enables higher fidelity qubit control independent of the qubit motional state, the latter being a source of error in laser-driven Raman gates. In both modes, the qubit states are read out by collecting state-dependent fluorescence on a photomultiplier tube array.

All superconducting-qubit experiments are executed on IBM's commercial Eagle and Heron processors, which host capacitively coupled transmon qubits arranged in a heavy-hex lattice and controlled with microwave single-qubit rotations, cross-resonance two-qubit gates, and dispersive resonator readout. We access these processors through the IBM quantum cloud application programming interface.

As an initial test of the BQNN, we classify 55 randomly selected MNIST test images using ten samples per image on QASM and IBM hardware. Results shown in Fig. 2(a) broadly agree with Ref. [5]. In particular, at intermediate quantumness ($a \approx 0.5$), the BQNN outperforms the analogous classical network. Experimental validation rates slightly exceed simulated values for $a \leq 0.5$. Although the difference lies within error bars, it can only be attributed to physical noise and suggests that moderate noise may

enhance performance. This behavior is analogous to stochastic noise injection in classical neural network training, which acts as an effective regularizer [31].

To further probe the impact of quantum-induced randomness and physical noise on image classification, we focus on a representative input image (index 6929) that is misclassified by the classical model ($a = 0$) but correctly classified in simulations once quantum noise is introduced ($a > 0$). We refer to such cases as No-Yes (NY) images. For comparison, images correctly classified by both the classical and quantum models are labeled Yes-Yes (YY) images. In simulation, outcome classes are defined by whether the classical model ($a = 0$) classifies an image correctly and whether, for $a > 0$, the quantum models achieve a classification accuracy of at least 50%. For each fixed model index a, we evaluate m independently trained models on N input images, repeating each evaluation r times to sample stochasticity arising from measurement. The resulting empirical frequencies of the outcome classes (YY, NY, YN, NN) are then computed. As N and r are increased, we expect these frequencies to converge to stable values, consistent with self-averaging behavior in the simulations. Across all simulations, we observe 77 NY and 9 YN outcomes, indicating a strong skew toward improved per-image performance in the presence of quantum effects. All results reported here are obtained for $a = 0$, 0.1, 0.2, 0.5, 1, with 100 samples per parameter value.

Perhaps the most interesting observation is the difference between the simulated behavior and experimental results at zero and small values of the quantumness parameter, a. At $a = 0$, the classical network deterministically fails and so does the idealized quantum simulation of its quantum realization. However, the experimental validation rate at $a = 0$ for this image is 10% for the trapped-ion quantum device and as high as 50% for the IBM device. This behavior too can only be attributed to physical noise. A plausible explanation is that there are two nearby minima in the energy landscape of the trained neural net [4]—one representing the erroneous reading, which gets selected classically 100% of the time, and the second being the correct reading. Measurement noise introduces classical stochasticity and errors, which in this case switch to the correct classification in a subset of runs. We note that we have not observed such sensitivity to physical noise at $a = 0$ for clear YY images (where presumably there is a single deep minimum insensitive to small perturbations). This simple observation suggests an interesting use case for noisy quantum neural networks as a way to flag ambiguous data entries. For the MNIST dataset, the data quality is relatively obvious from visual inspection, but this is not necessarily the case for more complex datasets.

We explore the validation rate as a function of quantumness, a, for the particular chosen image in Fig. 2(b). We observe a steep rise in the validation rate for this image to nearly 100% on the trapped-ion device and about 90% on the IBM quantum computer. When a becomes sufficiently large, it introduces nearly random rotations by producing qubit cat states with almost equal probabilities for the two computational basis states. This predictably hinders classification performance and we observe validation rates dropping with a. In the $a \to \infty$ limit classifications become uniformly random, and the validation rate converges to 0.1 in accordance with Eq. (3).

The validation rate inferred from both quantum platforms shows a subtle maximum at $a \approx 0.5$, but otherwise exhibits a plateau in an intermediate quantum regime before eventual performance decay. Owing to its sensitivity to physical noise, the $a = 0$ and optimal a regimes of the BQNN allow for a loose comparison of the physical noise between platforms. In particular, it is plausible that for these low-depth circuits, physical errors are the largest portion of the error. Hence, devices with relatively high measurement errors, such as superconducting devices, will see a greater noise-induced rise in validation rate [32]. In order to serve as a more stringent benchmark of quantum devices, entanglement and midcircuit measurement will need to be incorporated into the BQNN.

We further modify the BQNN to benchmark the effects of gate noise on feedforward dynamics. Specifically, n pairs of native unitary gates $UU^\dagger = I$ are injected into the circuit after each layer. While in an ideal system, they have no effect, their application in a real device enhances bit-flip, phase-flip, and other types of noise in the case of single-qubit pairs and introduce "parasitic entanglement" and crosstalk for pairs of time-reversed entangling gates.

Figure 2(c) shows results of our single-qubit noise experiments. In the classical regime, we observe for the trapped-ion device that even modest amounts of injected noise lead to a clear jump in NY-image accuracy at around 20 noise-gate pairs. We do not see any notable improvement of performance on the YY image with injected noise, which further validates the heuristic picture of two close-by minima for the borderline (ambiguous) NY data entries versus a single deep minimum for the clear YY images. This behavior carries over into the moderately quantum regime, where the combination of quantum randomness and added single-qubit noise has only a minor impact on validation rates. By contrast, the superconducting system does not exhibit any enhancement in performance on the NY images with injected noise, suggesting that the native physical measurement noise is already saturated or above the threshold to be beneficial for the BQNN performance. We see the per-image accuracy decline almost monotonically as noise increases.

The final test uses native two-qubit gate pairs and results are depicted in Fig. 3. For the trapped-ion experiment it was found that after the introduction of a *single* pair of two-qubit gates, the per-image classification rate shot up over

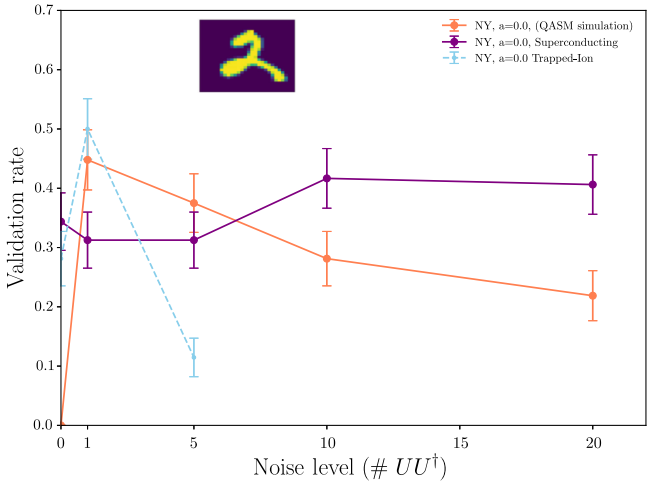


FIG. 3. For an NY image at $a = 0.0$, noise is controlled by the insertion of a number of two-qubit $UU^\dagger = I$ pairs.

40% before deteriorating to random 10% for five gate pairs. The superconducting platform was more resilient to this entanglement noise and even showed a small improvement in performance with the addition of further pairs, a behavior not consistent with our intuition or simulations.

Our results suggest that moderate physical noise, usually viewed as detrimental, can enhance certain quantum machine-learning tasks. While such effects at inference may be task-specific, they could provide more systematic advantages during training.

Extending this architecture toward genuinely classically nonsimulable quantum networks requires combining entangling gates, midcircuit measurements, and adaptive feedback. We recently developed the theory of such measurement-induced neural networks for a classically simulable subset of SU(4) gates [33], and efforts toward nonsimulable hardware implementations are ongoing.

Acknowledgments—This work is supported by a collaboration between the U.S. DOE and other agencies. This material is based upon work supported by the U.S. Department of Energy, Office of Science, National Quantum Information Science Research Centers, Quantum Systems Accelerator under Award No. DE-FOA-0002253 (X. L., A. M. G., and N. M. L.) and Basic Energy Sciences under Award No. DE-SC0001911 (V. G. and R. B.). Additional support is acknowledged from the Army Research Office under Grant No. W911NF-23-1-0241 (D. L.-H.) and from the National Science Foundation STAQ Project No. PHY-2325080 (N. M. L.).We acknowledge the use of IBM Quantum Credits for this work.

The views expressed are those of the authors, and do not reflect the official policy or position of IBM or the IBM Quantum team.

Data availability—The data supporting the findings of this study will be openly available in the GitHub repository at Galitski-theory-group/natural-quantization/data, with large files managed using Git Large File Storage (LFS), after May 27, 2026.

[1] J. J. Hopfield, Proc. Natl. Acad. Sci. U.S.A. **79**, 2554 (1982).
[2] D. J. Amit, H. Gutfreund, and H. Sompolinsky, Phys. Rev. A **32**, 1007 (1985).
[3] V. Dotsenko, *An Introduction to the Theory of Spin Glasses and Neural Networks*, World Scientific lecture notes in physics (World Scientific, Singapore, 1994), Vol. 55.
[4] R. Barney, M. Winer, and V. Galitski, arXiv:2408.06421.
[5] R. Barney, D. Lakhdar-Hamina, and V. Galitski, arXiv:2503.15482.
[6] K. L. Pudenz and D. A. Lidar, Quantum Inf. Process. **12**, 2027 (2013).
[7] S. Lloyd, M. Mohseni, and P. Rebentrost, arXiv:1307.0411.
[8] Z. Li, X. Liu, N. Xu, and J. Du, Phys. Rev. Lett. **114**, 140504 (2015).
[9] C. Ciliberto, M. Herbster, A. D. Ialongo, M. Pontil, A. Rocchetto, S. Severini, and L. Wossnig, Proc. Math. Phys. Eng. Sci. **474**, 20170551 (2018).
[10] M. Schuld, I. Sinayskiy, and F. Petruccione, Contemp. Phys. **56**, 172 (2015).
[11] J. Biamonte, P. Wittek, N. Pancotti, P. Rebentrost, N. Wiebe, and S. Lloyd, Nature (London) **549**, 195 (2017).
[12] Y. Wang and J. Liu, Rep. Prog. Phys. **87**, 116402 (2024).
[13] S. Lloyd, M. Schuld, A. Ijaz, J. Izaac, and N. Killoran, arXiv:2001.03622.
[14] J. Preskill, Quantum **2**, 79 (2018).
[15] A. Abbas, D. Sutter, C. Zoufal, A. Lucchi, A. Figalli, and S. Woerner, Nat. Comput. Sci. **1**, 403 (2021).
[16] H.-L. Huang *et al.*, Phys. Rev. Appl. **16**, 024051 (2021).
[17] X. Kong, L. Li, Z. Chen, C. Xue, X. Xu, H. Liu, Y. Wu, Y. Fang, H. Fang, K. Chen, Y. Yang, M. Dou, and G. Guo, arXiv:2503.12790.
[18] E. A. Cherrat, I. Kerenidis, N. Mathur, J. Landman, M. Strahm, and Y. Y. Li, Quantum **8**, 1265 (2024).
[19] P. Röseler, O. Schaudt, H. Berg, C. Bauckhage, and M. Koch, arXiv:2505.05957v2.
[20] P. Bermejo, P. Braccia, M. S. Rudolph, Z. Holmes, L. Cincio, and M. Cerezo, PRX Quantum **7**, 020304 (2026).
[21] E. Gil-Fuster, C. Gyurik, A. Perez-Salinas, and V. Dunjko, in *The Thirteenth International Conference on Learning Representations* (2025).
[22] Y. Li, X. Chen, and Matthew P. A. Fisher, Phys. Rev. B **100**, 134306 (2019).
[23] I. Goodfellow, Y. Bengio, and A. Courville, Gradient-based learning, in *Deep Learning* (MIT Press, Cambridge, MA, 2016), pp. 177–191.
[24] Y. Bengio, N. Léonard, and A. Courville, arXiv:1308.3432.
[25] National Institute of Standards and Technology, NIST special database 19: Handprinted forms and characters database, Available at https://www.nist.gov/srd/nist-special-database-19 (1990), accessed: 2025-06-12.
[26] A. Blais, R.-S. Huang, A. Wallraff, S. M. Girvin, and R. J. Schoelkopf, Phys. Rev. A **69**, 062320 (2004).

[27] J. Koch, T. M. Yu, J. Gambetta, A. A. Houck, D. I. Schuster, J. Majer, A. Blais, M. H. Devoret, S. M. Girvin, and R. J. Schoelkopf, Phys. Rev. A **76**, 042319 (2007).
[28] R. Barney, Natural quantization of neural networks, https://github.com/Galitski-theory-group/quantum-nn (2025).
[29] D. Lakhdar-Hamina, Natural-Quantization: A hybrid classical–quantum neural network framework, https://github.com/Galitski-theory-group/natural-quantization (2025).
[30] S. Debnath, N. M. Linke, C. Figgatt, K. A. Landsman, K. Wright, and C. Monroe, Nature (London) **536**, 63 (2016).
[31] C. M. Bishop, Neural Comput. **7**, 108 (1995).
[32] T. Tomesh, P. Gokhale, V. Omole, G. S. Ravi, K. N. Smith, J. Viszlai, X.-C. Wu, N. Hardavellas, M. R. Martonosi, and F. T. Chong, in *2022 IEEE International Symposium on High-Performance Computer Architecture (HPCA)* (IEEE Computer Society, Seoul, South Korea, 2022), pp. 587–603.
[33] P. Argyle, D. Lakhdar-Hamina, S. H. Miller, and V. Galitski, arXiv:2603.19200.