

Quantum generative adversarial learning in a superconducting quantum circuit

L. Hu,^{1,*} S. Wu,^{2,*} W. Cai,¹ Y. Ma,¹ X. Mu,¹ Y. Xu,¹ H. Wang,¹ Y. P. Song,¹ D.-L. Deng,^{1,†} C.-L. Zou,^{2,‡} and L. Sun^{1,§}

¹Center for Quantum Information, Institute for Interdisciplinary Information Sciences, Tsinghua University, Beijing 100084, China

²Key Laboratory of Quantum Information, CAS, University of Science and Technology of China, Hefei, Anhui 230026, P. R. China

Generative adversarial learning [1] is one of the most exciting recent breakthroughs in machine learning—a subfield of artificial intelligence that is currently driving a revolution in many aspects of modern society. It has shown splendid performance in a variety of challenging tasks such as image and video generations [2]. More recently, a quantum version of generative adversarial learning has been theoretically proposed [3, 4] and shown to possess the potential of exhibiting an exponential advantage over its classical counterpart [3]. Here, we report the first proof-of-principle experimental demonstration of quantum generative adversarial learning in a superconducting quantum circuit. We demonstrate that, after several rounds of adversarial learning, a quantum state generator can be trained to replicate the statistics of the quantum data output from a digital qubit channel simulator, with a high fidelity (98.8% on average) that the discriminator cannot distinguish between the true and the generated data. Our results pave the way for experimentally exploring the intriguing long-sought-after quantum advantages in machine learning tasks with noisy intermediate-scale quantum devices [5].

Machine learning [6], or more broadly artificial intelligence [7], represents an important area with general practical applications where near-term quantum devices may offer a significant speed-up over classical ones. With this vision, an intriguing interdisciplinary field of quantum machine learning/artificial intelligence has emerged and attracted tremendous attentions in recent years [8, 9]. A number of quantum algorithms that promise exponential speed-ups have been theoretically proposed [8–11] and some were demonstrated in proof-of-principle experiments [12, 13]. Yet, in most of these previous scenarios, the input data sets considered are typically classical. As a result, certain costly processes or techniques, such as quantum random access memories [14], are required to first map the classical data to quantum wave-vectors so as to be processed by quantum devices, rendering the potential speed-ups nullified [15].

In this Letter, we experimentally demonstrate a quantum version of generative adversarial network (QGAN) [3, 4], where both the input and output data sets are quantum from the beginning. In classical machine learning, a generative adversarial network (GAN) contains two major components, a generator (G) and a discriminator (D) [1]. They are trained through an adversarial learning procedure: in each learning round, D optimizes her strategies to identify the fake data produced by G, while G updates his strategies to fool D. Under reasonable assumptions, such an adversarial game will end up at a Nash equilibrium point, where G produces data that match the statistics of the true data, and D can no longer distinguish the fake data with a probability larger than $1/2$. In the quantum setting considered here, G consists of a superconducting circuit, which can generate an ensemble of quantum states with certain probability distribution, while D is composed of a quantum apparatus that carries out projective measurements. The arbitrary input quantum data is generated by a digital qubit channel simulator.

Figure 1(a) shows the schematic of the QGAN scheme. The black box provides the quantum true data which is described by a density matrix σ of a quantum system, while both the internal physical structure and the quantum process do not need

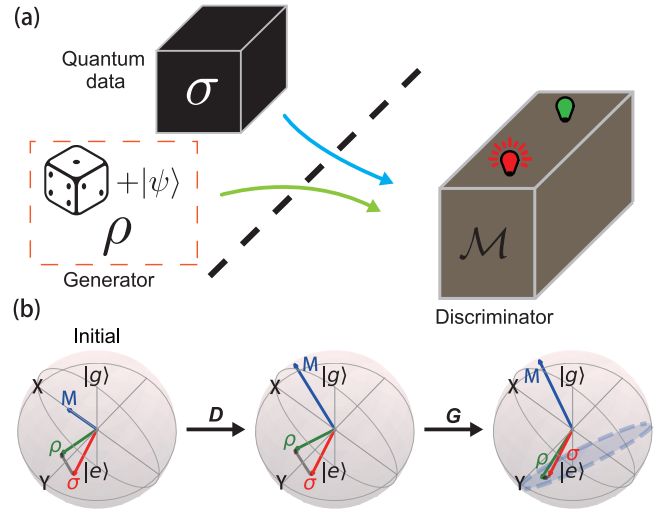


FIG. 1. **The quantum generative adversarial network (QGAN).** (a) The basic components of QGAN, including a black box quantum process for the quantum true data σ , the generator (G) that produces an ensemble of pure quantum states ρ , and the discriminator (D) that performs projective measurements \mathcal{M} . (b) The process of QGAN with the quantum states and the measurement basis on a Bloch sphere, where $\{|g\rangle, |e\rangle\}$ are the ground and excited state of a qubit. D and G play the adversarial game alternatively, in which D optimizes the measurement strategy to discriminate ρ from σ , while G optimizes the generation strategy to fool D.

to be known. G can generate arbitrary quantum states (ρ) by producing an ensemble of pure quantum states, i.e. a pure state from a set is randomly selected with certain probability to mimic the quantum true data. D performs quantum measurements (\mathcal{M}) on the true and the generated (fake) data, and attempts to distinguish them by the statistics of the measurement outcomes $p_\rho = \text{tr} \mathcal{M} \rho$ and $p_\sigma = \text{tr} \mathcal{M} \sigma$. In the QGAN, the measurement outcomes are public to both G and D. According to $p_{\rho, \sigma}$, D and G compete against each other by adaptively adjusting their strategies alternatively to distinguish ρ from σ and to fool D, respectively. It is interesting to note that σ and ρ represent two distinct interpretations of mixed quan-

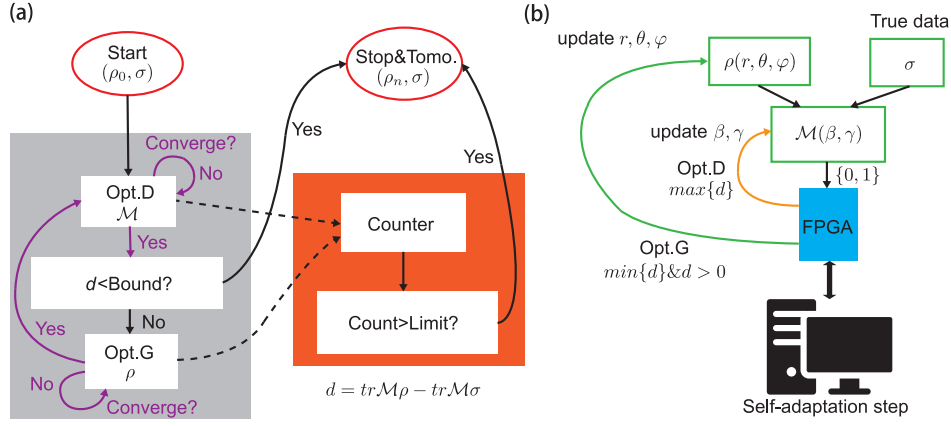


FIG. 2. **The experimental protocol of QGAN algorithm.** (a) The experiment starts with a quantum state σ as the quantum true data and a randomly generated state $\rho(r_0, \theta_0, \varphi_0)$ from G. Then, D and G optimize their strategies to compete against each other alternatively and repetitively. The stop condition of the game is either D fails to distinguish ρ from σ (the measurement output difference $d < d_B$, a preset bound) or the step count c_{step} reaches the limit c_B . (b) Procedure of optimizing D and G by the gradient decent method. The initial measurement axis $\mathcal{M}(\beta_0, \gamma_0)$ for D is randomly selected. The parameters β and γ are updated in the process of optimizing D, while r , θ , and φ are updated in the process of optimizing G. The measurement and control of the quantum system are realized through field programmable gate arrays (FPGA), while the estimations of the gradients are performed in a classical computer.

tum states: one is the output of a physical process in which an initial pure state might be entangled with some degrees of freedom of the environment; the other is an ensemble of pure states. Our QGAN scheme can also be explained as a game trying to distinguish between these two interpretations.

A visualized illustration of the general procedure of QGAN is depicted in Fig. 1(b) by presenting σ , ρ and \mathcal{M} of a qubit system in the Bloch sphere (note that we use the same notation \mathcal{M} to represent both the projective measurement and its corresponding axis). D and G play the game alternatively. D always starts first, and in her turn \mathcal{M} is optimized to maximize the difference of the measurement outcome $d = p_\rho - p_\sigma$. In an ideal case, D's turn ends up with $d = \frac{1}{2} \|\rho - \sigma\|_1$ corresponding to the normalized trace distance [16], and \mathcal{M} converges to be parallel with $\rho - \sigma$ in the Bloch sphere representation [Fig. 1(b)]. For G's turn, ρ is optimized to minimize d , and thus approaches a cross-section such that $\rho - \sigma$ is perpendicular to \mathcal{M} [Fig. 1(b)]. As a result, the trace distance between the fake and the true data reduces progressively in each round, and the game eventually approaches the unique Nash equilibrium with $d = 0$ and $p_\rho = p_\sigma = \frac{1}{2}$ [3].

We realize the QGAN learning algorithm [3] in a superconducting quantum electrodynamics architecture [17, 18]. Our experimental device consists of a superconducting transmon qubit dispersively coupled to a bosonic microwave mode [19, 20]. The quantum state of the transmon qubit serves as either ρ or σ alternatively in the algorithm. The bosonic mode facilitates the creation of the quantum true data σ in an arbitrary state through a digital quantum channel simulator, which requires adaptive control of both the transmon qubit and the bosonic mode. The detailed descriptions of the experimental device and apparatus are provided in the Supporting Materials, and can also be found in Ref. [21]. G generates the state $\rho(r, \theta, \varphi)$ of the transmon by randomly preparing a

pure state in the set $\{U(\theta, \varphi)|g\rangle, U(\pi - \theta, \varphi + \pi)|g\rangle\}$ with the corresponding probabilities $\{r, 1 - r\}$. Here, $U(\theta, \varphi) = e^{i\varphi\sigma_x/2}e^{i\theta\sigma_z/2}$ is the unitary operation on the transmon qubit, with σ_x and σ_z being the Pauli matrices. D performs the measurements by applying a unitary pre-rotating operation with the axis angles (β, γ) on the transmon and detecting the population of the ground state $|g\rangle$, which leads to $\mathcal{M} = U^\dagger(\beta, \gamma)|g\rangle\langle g|U(\beta, \gamma)$.

The protocol of our experimental QGAN algorithm is illustrated in Fig. 2(a). The experiment starts with a randomly generated state $\rho(r_0, \theta_0, \varphi_0)$ by G, a randomly picked measurement axis $\mathcal{M}(\beta_0, \gamma_0)$ by D, and the quantum true data σ from a fixed channel simulator. In each round of experiment, D plays the adversarial game first with ρ fixed, followed by G's turn with \mathcal{M} fixed. In all experiments, d is obtained by averaging $n = 5,000$ repetitive measurements on the true and the fake data respectively. The gradient $\partial d / \partial \xi$ for the control parameter $\xi \in \{r, \theta, \varphi, \beta, \gamma\}$ is critical for the QGAN. These gradients are approximately obtained by measuring d with respect to ξ and $\xi + \delta$ ($\delta \ll 1$) and calculating the differential numerically in a classical computer. According to the principle of gradient descent, the parameters are updated to maximize d (minimize d with $d > 0$) for D's (G's) turn, as explained in Fig. 2(b) (see Supporting Materials for the strategy). Here, the procedure to estimate each gradient for the relevant parameters is counted as one step, and the total number of steps quantifies the consumption of time and copies of data. In practical experiments, the projective detection outcomes follow a binomial statistic, and show a standard deviation of d as $\text{sd} = \sqrt{p_\rho(1 - p_\rho)/n + p_\sigma(1 - p_\sigma)/n}$. When approaching the Nash equilibrium, $p_\rho \approx p_\sigma \approx \frac{1}{2}$, then $\text{sd} \approx 1/\sqrt{2n} = 0.01$. Therefore, the measurement precision of d will limit the convergence of the game. In our experiments, D's turn ends when the differences of d in the last 3

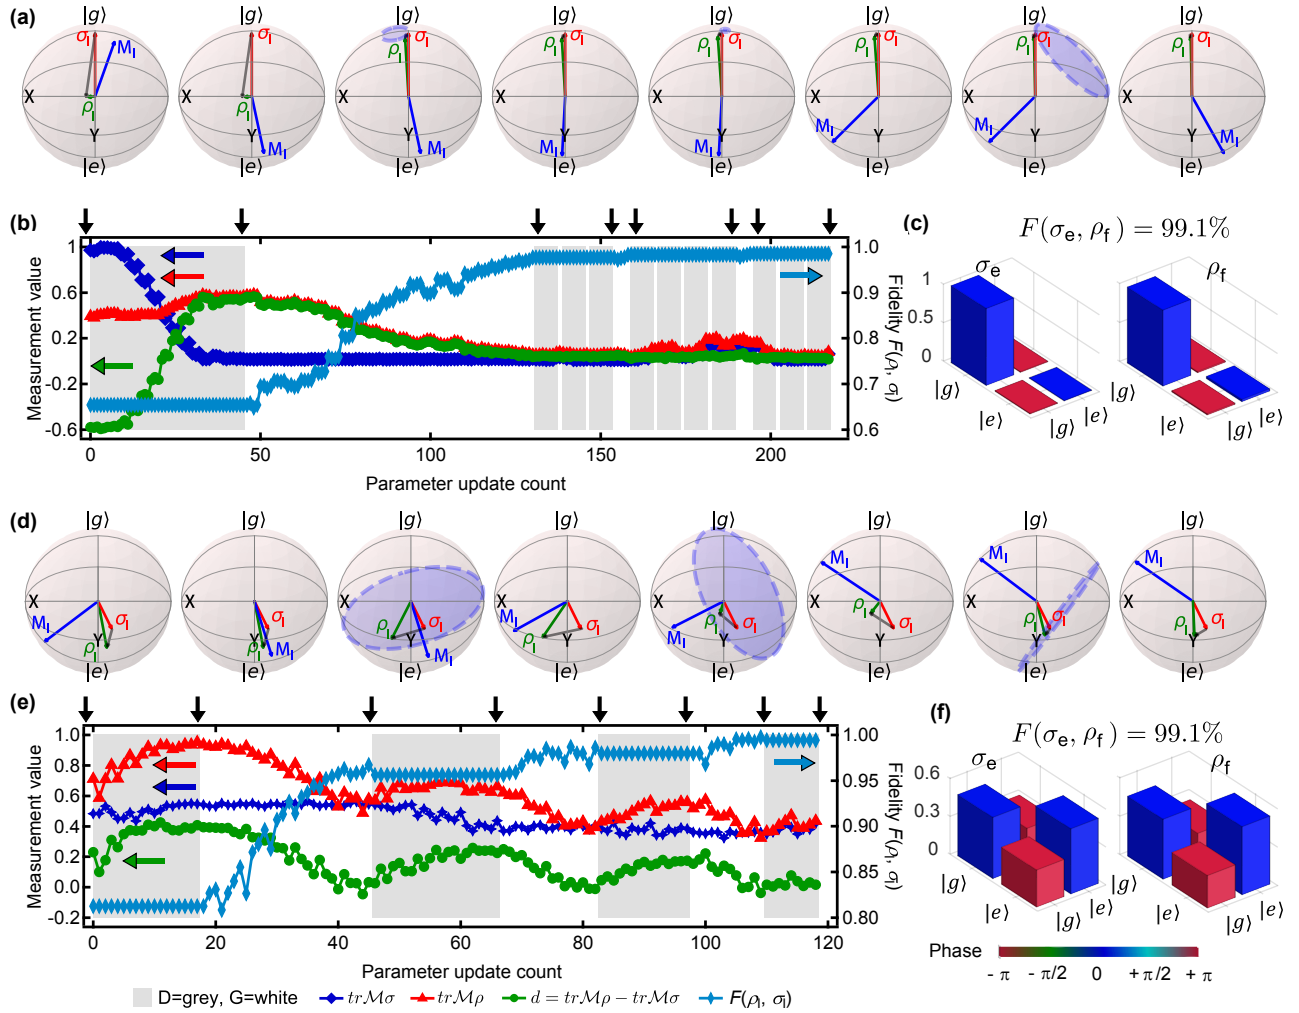


FIG. 3. **Tracking of QGAN.** (a-c) Experimental results for selecting $\sigma = |g\rangle\langle g|$ as the quantum true data. (a) shows the snapshots of the system at the particular steps indicated by arrows in (b) (from left to right in the same order) in the Bloch sphere representation. (b) shows the tracking of p_σ , p_ρ , d and F during the quantum adversarial learning process. The gray shadow regions are the processes of optimizing D, while the rests are those for optimizing G. (c) shows the measured state tomography of the experimental σ_e and final ρ_f with a state fidelity $F = 0.991$, demonstrating a successful QGAN that G can fool D by generating quantum data highly similar to the true data. (d-f) Typical experimental results for σ in an arbitrary mixed state with each panel being the counterpart of (a-c) respectively.

steps are less than 0.02. The G's turn ends when $d < R_j$ for the j th round: $R_j = 0.055 - 0.01j$ when $j \leq 3$ and $R_j = 0.02$ when $j > 3$. These two adversarial learning procedures can be repeated many rounds until either the total count of steps c_{step} reaches a preset limit c_B or the optimized d in D's round is smaller than a preset bound d_B .

Figures 3(a-c) show the typical results for the experimental QGAN with $\sigma = |g\rangle\langle g|$ of the transmon qubit, the highest purity state, as an example for the quantum true data. Since a digital quantum channel simulator can generate an arbitrary quantum state [21], the QGAN experiments by taking an arbitrary mixed state of the transmon as the true data is also studied and the results are depicted in Figs. 3(d-f). During the QGAN, the trajectory of control parameters are recorded [Figs. 3(b) and (e)] instead of characterizing the exact experimental ρ and \mathcal{M} . As shown in Figs. 3(a) and (d), the snapshots

of the quantum states and measurement axis at the particular steps indicated by the arrows in Figs. 3(b) and (e) respectively (from left to right in the same order) are plotted on the Bloch sphere. Here, σ_I , ρ_I and \mathcal{M}_I are the ideal results derived based on the calibrated control parameters. As expected, D adaptively adjusts \mathcal{M} to be parallel with $\rho_I - \sigma_I$, while G learns from the measurement outcomes to generate quantum data to fool D, and the generated quantum data gradually converges to the plane that contains the quantum true data and is perpendicular to \mathcal{M} . As a result, d oscillates in the D's and G's turns due to the adversarial process and eventually approaches 0, which indicates that ultimately D fails to discriminate ρ from σ , and G achieves his goal of replicating the statistics of the quantum true data.

To characterize the adversarial learning process, the state fidelity $F(\sigma_I, \rho_I) = \text{tr} \sqrt{\sqrt{\sigma_I} \rho_I \sqrt{\sigma_I}}$ in the adversarial pro-

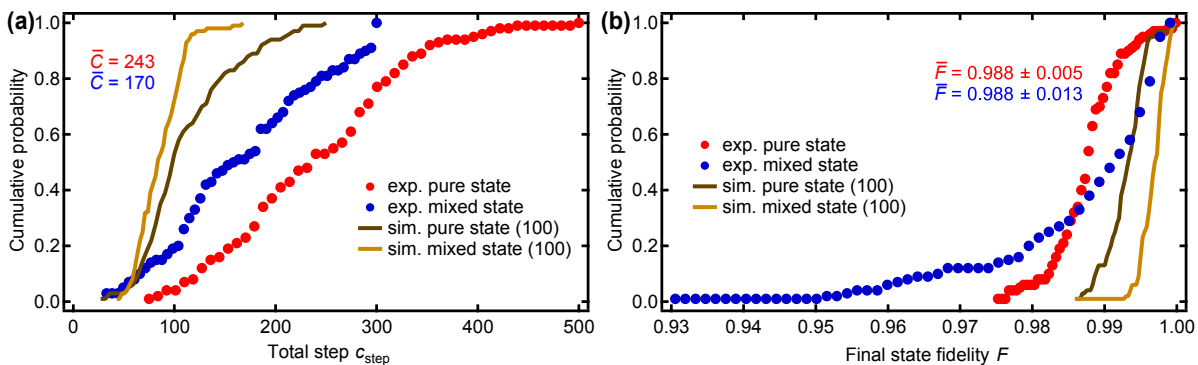


FIG. 4. **Statistics of QGAN performance.** (a) The cumulative probability of the total step count to finish the adversarial learning process. The QGAN is implemented for two difference cases, with a pure ($|g\rangle\langle g|$) and an arbitrary mixed state as the quantum true data, respectively. The count limit c_B for these two cases is 500 and 300, respectively. The obtained average c_{step} for these two types of adversarial learning process is 243 and 170 respectively. (b) The cumulative probability of final state fidelity F . The average fidelities for the pure and the mixed state are both 98.8%. For comparison, the noiseless numerical simulations of the adversarial learning process are also performed for 100 times, and their distributions are shown as solid lines.

cess is introduced to quantify the indistinguishability between the true data and the generated data. As plotted in Figs. 3(b) and (e), F approaches 1 after about 220 and 120 steps, respectively. The final generated quantum state ρ_f after the adversarial game and the experimental input state σ_e are measured by state tomography [Figs. 3(c) and (f)], and a fidelity as large as $F = 99.1\%$ for both cases are achieved. Such high fidelities verify that the unique Nash equilibrium, in which a quantum generator can replicate the statistics of the quantum true data, can be efficiently achieved in a quantum experimental realization of GAN. It is worth noting that, although we calibrate the system parameters to infer the ideal σ_I and ρ_I during the adversarial process, it is not necessary for the QGAN. Our experimental protocol can reach its equilibrium without requiring the knowledge about the exact data generated by G or the measurement axis chosen by D, and thus promises a double-blind quantum machine learning process just as its classical counterparts.

By taking the total steps (c_{step}) and the fidelity of final generated state (F) as the figures of merit, the statistics of our QGAN performance is finally studied with 100 random adversarial learning processes. We study both cases with the same pure and arbitrary mixed states as the quantum true data as in Fig. 3, but with different random ρ and \mathcal{M} at the beginning, all showing similar behaviors as those in Fig. 3. Figure 4(a) plots the cumulative probability of the total steps, i.e. the probability to finish the QGAN experiment within c_{step} steps. The average c_{step} for these two types of adversarial learning process are 243 and 170, respectively. Figure 4(b) shows the cumulative probability of state fidelity F with the average fidelities for both the pure and the mixed quantum data of 98.8%. Comparing to the noiseless numerical simulation results, the experimental average c_{step} is about twice larger, and the average F is about 1% lower. These differences are mainly attributed to the decoherence processes of the qubit, the finite measurement precision of d , and the non-ideal measured gradients. Further

studies about the effects of the experimental imperfections are provided in Supporting Materials.

In conclusion, we have demonstrated the feasibility of quantum generative adversarial learning with a superconducting quantum circuit, in which the input data, the generator and the discriminator are all quantum mechanical. Our results show that the generator can indeed learn the patterns of the input quantum data and produce quantum states with high fidelity that are not distinguishable by the discriminator. Since our QGAN experiment requires neither a quantum random accessing memory, nor a universal quantum computing device or any fine-tuning parameters (thus robust to experimental imperfections), it carries over to the noisy intermediate-scale quantum (NISQ) devices widely expected to be available in the near future [5]. An experimental demonstration of QGAN with NISQ devices promises to showcase the quantum advantages over classical GAN—a possible approach to realize quantum supremacy [22, 23] with practical applications. Our results may also have far-reaching consequences in solving quantum many-body problems with QGAN, given the recent rapid progresses in related directions [24–26]. In addition, the hybrid quantum-classical architecture demonstrated in this work can be straightforwardly extended to the optimal control [27] and self-guided quantum tomography [28], and we also anticipate their applications in other quantum machine learning/artificial intelligence algorithms.

* These two authors contributed equally to this work.

† dldeng@tsinghua.edu.cn

‡ clzou321@ustc.edu.cn

§ luyansun@tsinghua.edu.cn

[1] I. Goodfellow, J. Pouget-Abadie, M. Mirza, B. Xu, D. Warde-Farley, S. Ozair, A. Courville, and Y. Bengio, “Generative adversarial nets,” in *Advances in neural information processing*

- systems* (2014) pp. 2672–2680.
- [2] A. Creswell, T. White, V. Dumoulin, K. Arulkumaran, B. Sengupta, and A. A. Bharath, “Generative adversarial networks: An overview,” *IEEE Signal Processing Magazine* **35**, 53 (2018).
- [3] S. Lloyd and C. Weedbrook, “Quantum generative adversarial learning,” *Phys. Rev. Lett.* **121**, 40502 (2018).
- [4] P.-L. Dallaire-Demers and N. Killoran, “Quantum generative adversarial networks,” arXiv:1804.08641 (2018).
- [5] J. Preskill, “Quantum Computing in the NISQ era and beyond,” arXiv: 1801.00862 (2018).
- [6] M. Jordan and T. Mitchell, “Machine learning: Trends, perspectives, and prospects,” *Science* **349**, 255 (2015).
- [7] S. J. Russell and P. Norvig, *Artificial intelligence: a modern approach* (Malaysia; Pearson Education Limited., 2016).
- [8] J. Biamonte, P. Wittek, N. Pancotti, P. Rebentrost, N. Wiebe, and S. Lloyd, “Quantum machine learning,” *Nature* **549**, 195 (2017).
- [9] V. Dunjko and H. J. Briegel, “Machine learning & artificial intelligence in the quantum domain: a review of recent progress,” *Rep. Prog. Phys.* **81**, 074001 (2018).
- [10] C. Carlo, M. Herbster, A. Davide Ialongo, M. Pontil, A. Rocchetto, S. Severini, and L. Wossnig, “Quantum machine learning: a classical perspective,” *Proc.R.Soc.A* **474**, 20170551 (2017).
- [11] X. Gao, Z. Zhang, and L. Duan, “An efficient quantum algorithm for generative machine learning,” arXiv: 1711.02038 (2017).
- [12] J. S. Otterbach, R. Manenti, N. Alidoust, A. Bestwick, M. Block, B. Bloom, S. Caldwell, N. Didier, E. S. Fried, S. Hong, P. Karalekas, C. B. Osborn, A. Papageorge, E. C. Peterson, G. Prawiroatmodjo, N. Rubin, C. A. Ryan, D. Scarbelli, M. Scheer, E. A. Sete, P. Sivarajah, R. S. Smith, A. Staley, N. Tezak, W. J. Zeng, A. Hudson, B. R. Johnson, M. Reagor, M. P. da Silva, and C. Rigetti, “Unsupervised Machine Learning on a Hybrid Quantum Computer,” arXiv: 1712.05771 (2017).
- [13] Z. Li, X. Liu, N. Xu, and J. Du, “Experimental realization of a quantum support vector machine,” *Phys. Rev. Lett.* **114**, 140504 (2015).
- [14] V. Giovannetti, S. Lloyd, and L. Maccone, “Quantum Random Access Memory,” *Phys. Rev. Lett.* **100**, 160501 (2008).
- [15] A. Scott, “Read the fine print,” *Nat. Phys.* **11**, 291 (2015).
- [16] M. A. Nielsen and I. L. Chuang, *Quantum Computation and Quantum Information* (Cambridge Univ. Press, 2000).
- [17] M. H. Devoret and R. J. Schoelkopf, “Superconducting circuits for quantum information: an outlook.” *Science* **339**, 1169 (2013).
- [18] X. Gu, A. F. Kockum, A. Miranowicz, Y. xi Liu, and F. Nori, “Microwave photonics with superconducting quantum circuits,” *Phys. Rep.* **718-719**, 1 (2017).
- [19] A. Wallraff, D. I. Schuster, A. Blais, L. Frunzio, R.-S. Huang, J. Majer, S. Kumar, S. M. Girvin, and R. J. Schoelkopf, “Circuit quantum electrodynamics: Coherent coupling of a single photon to a cooper pair box,” *Nature* **431**, 162 (2004).
- [20] H. Paik, D. I. Schuster, L. S. Bishop, G. Kirchmair, G. Catelani, A. P. Sears, B. R. Johnson, M. J. Reagor, L. Frunzio, L. I. Glazman, S. M. Girvin, M. H. Devoret, and R. J. Schoelkopf, “Observation of High Coherence in Josephson Junction Qubits Measured in a Three-Dimensional Circuit QED Architecture,” *Phys. Rev. Lett.* **107**, 240501 (2011).
- [21] L. Hu, X. Mu, W. Cai, Y. Ma, Y. Xu, H. Wang, Y. Song, C.-L. Zou, and L. Sun, “Experimental repetitive quantum channel simulation,” arXiv , 1807.07694 (2018).
- [22] J. Preskill, “Quantum computing and the entanglement frontier,” arXiv:1203.5813 (2012).
- [23] A. W. Harrow and A. Montanaro, “Quantum computational supremacy,” *Nature* **549**, 203 (2017).
- [24] G. Carleo and M. Troyer, “Solving the quantum many-body problem with artificial neural networks,” *Science* **355**, 602 (2017).
- [25] D.-L. Deng, X. Li, and S. Das Sarma, “Quantum entanglement in neural network states,” *Phys. Rev. X* **7**, 021021 (2017).
- [26] D.-L. Deng, “Machine learning detection of bell nonlocality in quantum many-body systems,” *Phys. Rev. Lett.* **120**, 240402 (2018).
- [27] J. Li, X. Yang, X. Peng, and C.-P. Sun, “Hybrid Quantum-Classical Approach to Quantum Optimal Control,” *Phys. Rev. Lett.* **118**, 150503 (2017).
- [28] R. J. Chapman, C. Ferrie, and A. Peruzzo, “Experimental Demonstration of Self-Guided Quantum Tomography,” *Phys. Rev. Lett.* **117**, 040402 (2016).

Acknowledgments

We thank N. Ofek and Y. Liu for valuable suggestions on FPGA programming. LS acknowledges the support from National Key Research and Development Program of China No.2017YFA0304303, National Natural Science Foundation of China Grant No.11474177, and the Thousand Youth Fellowship program in China. LS also thanks R. Vijay and his group for help on the parametric amplifier measurements. CLZ is supported by Anhui Initiative in Quantum Information Technologies (AHY130000). D.L.D. acknowledges the start-up fund from Tsinghua University.