

PHYSICS

Quantum generative adversarial learning in a superconducting quantum circuit

Ling Hu^{1*}, Shu-Hao Wu^{2*}, Weizhou Cai¹, Yuwei Ma¹, Xianghao Mu¹, Yuan Xu¹, Haiyan Wang¹, Yipu Song¹, Dong-Ling Deng^{1†}, Chang-Ling Zou^{2†}, Luyan Sun^{1†}

Generative adversarial learning is one of the most exciting recent breakthroughs in machine learning. It has shown splendid performance in a variety of challenging tasks such as image and video generation. More recently, a quantum version of generative adversarial learning has been theoretically proposed and shown to have the potential of exhibiting an exponential advantage over its classical counterpart. 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 quantum channel simulator, with a high fidelity (98.8% on average) so 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.

INTRODUCTION

Machine learning (1), or more broadly artificial intelligence (2), represents an important area with general practical applications where near-term quantum devices may offer a substantial speedup over classical ones. With this vision, an intriguing interdisciplinary field of quantum machine learning/artificial intelligence has emerged and attracted tremendous attention in recent years (3, 4). A number of quantum algorithms that promise exponential speedups have been theoretically proposed (3–6), and some were demonstrated in proof-of-principle experiments (7, 8). Yet, in most of these previous scenarios, the input datasets considered are typically classical. As a result, certain costly processes or techniques, such as quantum random access memories (9), are required to first map the classical data to quantum wave functions so as to be processed by quantum devices, rendering the potential speedups nullified (10).

Here, we experimentally demonstrate a quantum version of generative adversarial network (QGAN) (11, 12), where both input and output datasets are quantum from the beginning. In classical machine learning, a GAN (13, 14) contains two major components, a generator (G) and a discriminator (D) (13). 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 are generated by a quantum channel simulator.

RESULTS

The QGAN algorithm

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




Fig. 1. The QGAN. (A) The basic components of the 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 the 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.

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, China.

*These authors contributed equally to this work.

†Corresponding author. Email: dldeng@tsinghua.edu.cn (D.-L.D.); czou321@ustc.edu.cn (C.-L.Z.); luyansun@tsinghua.edu.cn (L.S.)




Fig. 2. The experimental protocol of the QGAN algorithm. (A) The experiment starts with a state σ as the quantum true data and a randomly generated state $\rho(r_0, \theta_0, \varphi_0)$ from the generator. Then, the discriminator (D) and generator (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 using the gradient descent method. The initial measurement axis $\mathcal{M}(\beta_0, \gamma_0)$ for D is randomly chosen. 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 (FPGAs), while the estimations of the gradients are performed on a classical computer.

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. σ and ρ are two distinct interpretations of mixed quantum 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 the QGAN is depicted in Fig. 1B 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 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 $d = \frac{1}{2} \|\rho - \sigma\|$, corresponding to the normalized trace distance (15), and \mathcal{M} converges to be parallel with $\rho - \sigma$ in the Bloch sphere representation (Fig. 1B). 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. 1B). 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 $p_{\sigma} = p_{\rho} = \frac{1}{2}$ (11).

Experimental implementation

We realize the QGAN learning algorithm (11) in a superconducting quantum electrodynamics architecture (16, 17). Our experimental device consists of a superconducting transmon qubit dispersively coupled to a bosonic microwave mode (18–23). 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 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 Supplementary Materials and (24, 25). 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, $U(\theta, \varphi) = e^{i\varphi\sigma_z/2} e^{i\theta\sigma_x/2}$ 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 prerotating 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. 2A. 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 quantum 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 runs, d is obtained by averaging $n = 5000$ 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 on 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. 2B (see the Supplementary Materials for the strategy). Here, each determination of d 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 (sd) of d as $\text{sd} = \sqrt{p_{\rho}(1 - p_{\rho})/n + p_{\sigma}(1 - p_{\sigma})/n}$. When approaching the Nash equilibrium, $p_{\rho} \approx p_{\sigma} \approx \frac{1}{2}$, and 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 three 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 .

The QGAN performance

Figure 3 (A to C) shows the typical results for the experimental QGAN with $\sigma = |g\rangle\langle g|$ of the transmon qubit, the highest purity state that can be achieved in the experiment, as an example for the quantum true data. Since a quantum channel simulator can generate an arbitrary quantum state (25), the QGAN experiments by taking an arbitrary mixed state of




Fig. 3. Tracking of the QGAN. (A to C) Experimental results for selecting $\sigma = |g\rangle|g\rangle$ as the quantum true data. (A) The snapshots of the system at the particular steps indicated by the black vertical arrows in (B) (from left to right in the same order), in the Bloch sphere representation. (B) 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. The horizontal color arrows indicate the vertical axis that each curve with the same color corresponds to. Since the convergence condition d_b for the case of pure states is small (see the Supplementary Materials) and there is inevitable measurement imprecision, the optimized \mathcal{M} is difficult to obtain or could even be randomized in certain trials. In this particular trial, \mathcal{M} ends up nearly antiparallel with both ρ and σ , resulting in $\text{tr} \mathcal{M} \sigma \approx \text{tr} \mathcal{M} \rho \approx 0$ and $d \approx 0$. (C) 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 to F) Typical experimental results for σ in an arbitrary mixed state with each panel being the counterpart of (A) to (C), respectively.

the transmon as the true data is also studied, and the results are depicted in Fig. 3 (D to F). During the QGAN, the trajectory of control parameters are recorded (Fig. 3, B and E) instead of characterizing the exact experimental ρ and \mathcal{M} . As shown in Fig. 3 (A and D), the snapshots of the quantum states and measurement axis at the particular steps indicated by the arrows in Fig. 3 (B and E, respectively; from left to right in the same order) are plotted on the Bloch sphere. Here, σ_T, ρ_T , and \mathcal{M}_T are the ideal results derived on the basis of the calibrated control parameters. As expected, D adaptively adjusts \mathcal{M} to be parallel with $\rho_T - \sigma_T$, 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 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, we introduce the state fidelity $F(\sigma_T, \rho_T) = \text{tr} \sqrt{\sqrt{\sigma_T} \rho_T \sqrt{\sigma_T}}$ in the adversarial process to quantify the indistinguishability between the true and the generated data. As plotted in Fig. 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 using state tomography (Fig. 3, C and F), and fidelities as large as $F(\sigma_e, \rho_f) = 99.1\%$ are achieved for both the pure- and mixed-state cases. These high fidelities verify that the unique Nash equilibrium, in which a quantum G can replicate the statistics of the quantum true data, can be efficiently achieved in a quantum experimental realization of GAN. Although we calibrate the system parameters to infer the ideal σ_T and ρ_T

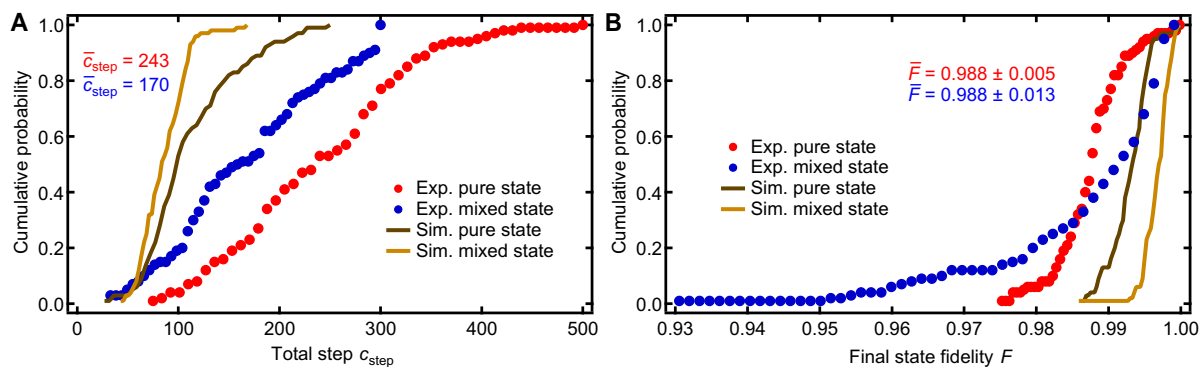


Fig. 4. Statistics of the QGAN performance. (A) The cumulative probability of the total step count to finish the adversarial learning process. The QGAN is implemented for two different cases, with a pure ($|g\rangle|g\rangle$) and an arbitrary mixed state as the quantum true data, respectively. The count 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. Exp., experimental; sim., simulation. (B) The cumulative probability of final state fidelity F . The average fidelities for the pure state and the mixed state are both 98.8%. For comparison, the noiseless numerical simulations of the adversarial learning process are also performed 100 times, and their distributions are shown as solid lines.

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 the final generated state (F) as the figures of merit, the statistics of our QGAN performance is lastly 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 4A 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 4B 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 the Supplementary Materials.

DISCUSSION

The QGAN algorithm demonstrated in our experiments can be directly extended to a quantum system with higher dimensions. In our superconducting architecture, the bosonic mode actually provides a quantum system with infinite dimensions, which can be encoded as a multilevel system. On the basis of the same adaptive technique used in our current experiment, an arbitrary quantum state of a photonic qudit with m levels (equivalent to a $\log_2 m$ -qubit system) can be generated and manipulated with the assistance of the transmon qubit (25). Then, our experiment can be straightforwardly extended to this photonic qudit that serves as either the quantum true or generated data. Another possible extension of our current experiment is to explore a more complicated architecture with multiple bosonic modes coupled to multiple transmon qubits (26–28). Both the quantum true and generated data can be stored in these bosonic modes.

For an m -dimensional system, the number of the QGAN parameters scales as $O(m^2)$; therefore, our algorithm is still feasible for experiments with reasonable time and quantum resource consumptions. For instance, our numerical simulations (see the Supplementary Materials) indicate that the QGAN algorithm for two- and three-qubit systems could converge to a final state fidelity larger than 0.95 with roughly thousands of steps on average. In the current experiment, the gradient is estimated numerically via a classical mean. As envisioned in (12), the gradient can also be obtained through a quantum circuit with higher efficiency if more quantum resources are available. We note that, although quantum-state tomography is applied in our experiment to characterize the performance, it is not necessary for the QGAN algorithm. In the case of a quantum system with a large number of qubits, a better and more efficient way of assessing the QGAN performance would be desired.

In conclusion, we have demonstrated the feasibility of quantum generative adversarial learning with a superconducting quantum circuit in which the input data, the G, and the D are all quantum mechanical. Our results show that the G can learn the patterns of the input quantum data and produce quantum states with high fidelity that are not distinguishable by the D. 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 (29) widely expected to be available in the near future. An experimental demonstration of the QGAN with NISQ devices promises to showcase the quantum advantages over classical GAN—a possible approach to realizing quantum supremacy (30, 31) with practical applications. Yet, an unambiguous experimental demonstration of quantum supremacy in this context may still require substantial technology advancements. Our results might also have far-reaching consequences in solving quantum many-body problems with the QGAN algorithm, given the recent rapid progress in related directions (32–36). In addition, the hybrid quantum-classical architecture demonstrated in this work can be straightforwardly extended to the optimal control (37) and self-guided quantum tomography (38), and we also anticipate their applications in other quantum machine learning/artificial intelligence algorithms.

MATERIALS AND METHODS

Our experimental device consists of a superconducting transmon qubit dispersively coupled to two microwave cavities (18, 19). The transmon qubit has an energy relaxation time $T_1 = 30 \mu\text{s}$ and a pure dephasing time $T_\phi = 120 \mu\text{s}$, and its quantum state serves as either ρ or σ alternatively during the QGAN algorithm. The long-lived cavity has a photon lifetime $T_1 = 143 \mu\text{s}$ and serves as an ancillary photonic qubit to facilitate the creation of the quantum true data σ in an arbitrary state through a quantum channel simulator (25). In (25), we showed that an arbitrary quantum state of the photonic qubit can be generated deterministically, which can be mapped back onto the transmon qubit for the initial σ . The other short-lived cavity with $T_1 = 44 \text{ ns}$ is to readout the transmon qubit with the help of a phase-sensitive Josephson bifurcation amplifier (39–41) for a high-fidelity single-shot measurement. In each experiment, the qubit is initialized to the ground state $|g\rangle$. The operations on the transmon qubit and the ancillary photonic qubit are realized through numerically optimized pulse sequences with the gradient ascent pulse engineering method (42, 43) based on carefully calibrated system parameters and are implemented by field programmable gate arrays (FPGAs) with customized logic. The details of the FPGAs and the experimental apparatus can be found in 24.

In our experiments, one step consists of $2n = 10,000$ measurements of the quantum states (both ρ and σ) for calculating d , which is selected as the loss function for the QGAN optimization. During the turns of D and G, the gradients of d were estimated with respect to each optimization parameter. Each update of the parameters requires $n_p + 1$ steps for n_p parameters, with $n_p = 2$ and 3 for D and G, respectively. State tomography was performed after the QGAN algorithm to characterize the true data σ and the generated ρ based on the recorded parameters in the quantum channel simulator and the QGAN algorithm. Note that we did not repeat the QGAN process to obtain the tomography to save time.

SUPPLEMENTARY MATERIALS

Supplementary material for this article is available at <http://advances.sciencemag.org/cgi/content/full/5/1/eaav2761/DC1>

Section S1. Tracking of $\rho - \sigma$ and \mathcal{M} during the QGAN process

Section S2. Comparisons between the experimental and the numerical simulation results

Section S3. Algorithm and numerical results

Fig. S1. Tracking of $\rho - \sigma$ and \mathcal{M} and the comparison between experiments and numerical simulations based on the recorded parameters ξ in the QGAN process as shown in Fig. 3.

Fig. S2. Statistics of the QGAN performance based on numerical simulations.

Fig. S3. Influence of δ and sd in the QGAN performance and count of steps when the QGAN converges to the same d_b for different γ and p .

Fig. S4. Performance of the QGAN algorithm for multipartite quantum systems with the quantum true data being a random mixed state.

Fig. S5. Performance of the QGAN algorithm with Greenberger-Horne-Zeilinger and W states as the quantum true data.

REFERENCES AND NOTES

- M. I. Jordan, T. M. Mitchell, Machine learning: Trends, perspectives, and prospects. *Science* **349**, 255–260 (2015).
- S. J. Russell, P. Norvig, *Artificial Intelligence: A Modern Approach* (Pearson Education Limited, 2016).
- J. Biamonte, P. Wittek, N. Pancotti, P. Rebentrost, N. Wiebe, S. Lloyd, Quantum machine learning. *Nature* **549**, 195–202 (2017).
- V. Dunjko, H. J. Briegel, Machine learning & artificial intelligence in the quantum domain: A review of recent progress. *Rep. Prog. Phys.* **81**, 074001 (2018).
- C. Ciliberto, M. Herbster, A. D. Ialongo, M. Pontil, A. Rocchetto, S. Severini, L. Wossnig, Quantum machine learning: A classical perspective. *Proc. R. Soc. A* **474**, 20170551 (2018).
- X. Gao, Z. Zhang, L. Duan, A quantum machine learning algorithm based on generative models. *Sci. Adv.* **4**, eaat9004 (2019).
- 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. Scarabelli, 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, C. Rigetti, Unsupervised machine learning on a hybrid quantum computer. arXiv:1712.05771 [quant-ph] (15 December 2017).
- Z. Li, X. Liu, N. Xu, J. Du, Experimental realization of a quantum support vector machine. *Phys. Rev. Lett.* **114**, 140504 (2015).
- V. Giovannetti, S. Lloyd, L. Maccone, Quantum random access memory. *Phys. Rev. Lett.* **100**, 160501 (2008).
- S. Aaronson, Read the fine print. *Nat. Phys.* **11**, 291–293 (2015).
- S. Lloyd, C. Weedbrook, Quantum generative adversarial learning. *Phys. Rev. Lett.* **121**, 040502 (2018).
- P.-L. Dallaire-Demers, N. Killoran, Quantum generative adversarial networks. *Phys. Rev. A* **98**, 012324 (2018).
- I. J. Goodfellow, J. Pouget-Abadie, M. Mirza, B. Xu, D. Warde-Farley, S. Ozair, A. Courville, Y. Bengio, Generative adversarial nets, in *Advances in Neural Information Processing Systems* (MIT Press, 2014), pp. 2672–2680.
- A. Creswell, T. White, V. Dumoulin, K. Arulkumar, B. Sengupta, A. A. Bharath, Generative adversarial networks: An overview. *IEEE Signal Process. Mag.* **35**, 53–65 (2018).
- M. A. Nielsen, I. L. Chuang, *Quantum Computation and Quantum Information* (Cambridge Univ. Press, 2000).
- M. H. Devoret, R. J. Schoelkopf, Superconducting circuits for quantum information: An outlook. *Science* **339**, 1169–1174 (2013).
- X. Gu, A. F. Kockum, A. Miranowicz, Y.-x. Liu, F. Nori, Microwave photonics with superconducting quantum circuits. *Phys. Rep.* **718–719**, 1–102 (2017).
- A. Wallraff, D. I. Schuster, A. Blais, L. Frunzio, R.-S. Huang, J. Majer, S. Kumar, S. M. Girvin, R. J. Schoelkopf, Circuit quantum electrodynamics: Coherent coupling of a single photon to a cooper pair box. *Nature* **431**, 162–167 (2004).
- 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, 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).
- G. Kirchmair, B. Vlastakis, Z. Leghtas, S. E. Nigg, H. Paik, E. Ginossar, M. Mirrahimi, L. Frunzio, S. M. Girvin, R. J. Schoelkopf, Observation of quantum state collapse and revival due to the single-photon Kerr effect. *Nature* **495**, 205–209 (2013).
- D. Ristè, M. Dukalski, C. A. Watson, G. de Lange, M. J. Tiggelman, Y. M. Blanter, K. W. Lehnert, R. N. Schouten, L. DiCarlo, Deterministic entanglement of superconducting qubits by parity measurement and feedback. *Nature* **502**, 350–354 (2013).
- B. Vlastakis, G. Kirchmair, Z. Leghtas, S. E. Nigg, L. Frunzio, S. M. Girvin, M. Mirrahimi, M. H. Devoret, R. J. Schoelkopf, Deterministically encoding quantum information using 100-photon Schrödinger cat states. *Science* **342**, 607–610 (2013).
- N. Ofek, A. Petrenko, R. Heeres, P. Reinhold, Z. Leghtas, B. Vlastakis, Y. Liu, L. Frunzio, S. M. Girvin, L. Jiang, M. Mirrahimi, M. H. Devoret, R. J. Schoelkopf, Extending the lifetime of a quantum bit with error correction in superconducting circuits. *Nature* **536**, 441–445 (2016).
- L. Hu, Y. Ma, W. Cai, X. Mu, Y. Xu, W. Wang, Y. Wu, H. Wang, Y. Song, C. Zou, S. M. Girvin, L.-M. Duan, L. Sun, Demonstration of quantum error correction and universal gate set on a binomial bosonic logical qubit. arXiv:1805.09072 [quant-ph] (23 May 2018).
- L. Hu, X. Mu, W. Cai, Y. Ma, Y. Xu, H. Wang, Y. Song, C.-L. Zou, L. Sun, Experimental repetitive quantum channel simulation. *Sci. Bull.* **63**, 1551–1557 (2018).
- K. S. Chou, J. Z. Blumoff, C. S. Wang, P. C. Reinhold, C. J. Axline, Y. Y. Gao, L. Frunzio, M. H. Devoret, L. Jiang, R. J. Schoelkopf, Deterministic teleportation of a quantum gate between two logical qubits. *Nature* **561**, 368–373 (2018).
- Y. Y. Gao, B. J. Lester, Y. Zhang, C. Wang, S. Rosenblum, L. Frunzio, L. Jiang, S. M. Girvin, R. J. Schoelkopf, Programmable interference between two microwave quantum memories. *Phys. Rev. X* **8**, 021073 (2018).
- Y. Xu, W. Cai, Y. Ma, X. Mu, W. Dai, W. Wang, L. Hu, X. Li, J. Han, H. Wang, Y. Song, Z.-B. Yang, S.-B. Zheng, L. Sun, Geometrically manipulating photonic Schrödinger cat states and realizing cavity phase gates. arXiv:1810.04690 [quant-ph] (10 October 2018).
- J. Preskill, Quantum computing in the NISQ era and beyond. arXiv:1801.00862 [quant-ph] (2 January 2018).
- J. Preskill, Quantum computing and the entanglement frontier; arXiv:1203.5813 [quant-ph] (26 March 2012).
- A. W. Harrow, A. Montanaro, Quantum computational supremacy. *Nature* **549**, 203–209 (2017).
- G. Carleo, M. Troyer, Solving the quantum many-body problem with artificial neural networks. *Science* **355**, 602–606 (2017).
- L. Wang, Discovering phase transitions with unsupervised learning. *Phys. Rev. B* **94**, 195105 (2016).
- E. P. L. van Nieuwenburg, Y.-H. Liu, S. D. Huber, Learning phase transitions by confusion. *Nat. Phys.* **13**, 435–439 (2017).

35. J. Carrasquilla, R. G. Melko, Machine learning phases of matter. *Nat. Phys.* **13**, 431–434 (2017).
36. D.-L. Deng, X. Li, S. Das Sarma, Quantum entanglement in neural network states. *Phys. Rev. X* **7**, 021021 (2017).
37. J. Li, X. Yang, X. Peng, C.-P. Sun, Hybrid quantum-classical approach to quantum optimal control. *Phys. Rev. Lett.* **118**, 150503 (2017).
38. R. J. Chapman, C. Ferrie, A. Peruzzo, Experimental demonstration of self-guided quantum tomography. *Phys. Rev. Lett.* **117**, 040402 (2016).
39. M. Hatridge, R. Vijay, D. H. Slichter, J. Clarke, I. Siddiqi, Dispersive magnetometry with a quantum limited SQUID parametric amplifier. *Phys. Rev. B* **83**, 134501 (2011).
40. T. Roy, S. Kundu, M. Chand, A. M. Vadiraj, A. Ranadive, N. Nehra, M. P. Patankar, J. Aumentado, A. A. Clerk, R. Vijay, Broadband parametric amplification with impedance engineering: Beyond the gain-bandwidth product. *Appl. Phys. Lett.* **107**, 262601 (2015).
41. A. Kamal, A. Marblestone, M. Devoret, Signal-to-pumpback action and self-oscillation in double-pump Joseph-son parametric amplifier. *Phys. Rev. B* **79**, 184301 (2009).
42. N. Khaneja, T. Reiss, C. Kehlet, T. Schulte-Herbrüggen, S. J. Glaser, Optimal control of coupled spin dynamics: Design of NMR pulse sequences by gradient ascent algorithms. *J. Magn. Reson.* **172**, 296–305 (2005).
43. P. de Fouquieres, S. Schirmer, S. Glaser, I. Kuprov, Second order gradient ascent pulse engineering. *J. Magn. Reson.* **212**, 412–417 (2011).

Acknowledgments: We thank N. Ofek and Y. Liu for valuable suggestions on FPGA programming. L.S. also thanks R. Vijay and co-workers for help on the parametric amplifier

measurements. **Funding:** L.S. acknowledges the support from National Key Research and Development Program of China grant no. 2017YFA0304303 and National Natural Science Foundation of China grant no. 11474177. C.-L.Z. is supported by Anhui Initiative in Quantum Information Technologies (AHY130000). D.-L.D. acknowledges the start-up fund from Tsinghua University. **Author contributions:** L.H. performed the experiment and analyzed the data with the support from W.C. and Y.M. S.-H.W., D.-L.D., and C.-L.Z. provided theoretical supports. L.H. and X.M. fabricated devices with the assistance of Y.X., H.W., and Y.S. C.-L.Z. proposed the experiment. L.S. supervised the experiment. D.-L.D., C.-L.Z., and L.S. wrote the manuscript with feedback from all authors. **Competing interests:** The authors declare that they have no competing interests. **Data and materials availability:** All data needed to evaluate the conclusions in the paper are present in the paper and/or the Supplementary Materials. Additional data related to this paper may be requested from the authors.

Submitted 31 August 2018
Accepted 11 December 2018
Published 25 January 2019
10.1126/sciadv.aav2761

Citation: L. Hu, S.-H. Wu, W. Cai, Y. Ma, X. Mu, Y. Xu, H. Wang, Y. Song, D.-L. Deng, C.-L. Zou, L. Sun, Quantum generative adversarial learning in a superconducting quantum circuit. *Sci. Adv.* **5**, eaav2761 (2019).

Quantum generative adversarial learning in a superconducting quantum circuit

Ling Hu, Shu-Hao Wu, Weizhou Cai, Yuwei Ma, Xianghao Mu, Yuan Xu, Haiyan Wang, Yipu Song, Dong-Ling Deng, Chang-Ling Zou and Luyan Sun

Sci Adv 5 (1), eaav2761.
DOI: 10.1126/sciadv.aav2761

ARTICLE TOOLS

<http://advances.sciencemag.org/content/5/1/eaav2761>

SUPPLEMENTARY MATERIALS

<http://advances.sciencemag.org/content/suppl/2019/01/18/5.1.eaav2761.DC1>

REFERENCES

This article cites 35 articles, 4 of which you can access for free
<http://advances.sciencemag.org/content/5/1/eaav2761#BIBL>

PERMISSIONS

<http://www.sciencemag.org/help/reprints-and-permissions>

Use of this article is subject to the [Terms of Service](#)

Science Advances (ISSN 2375-2548) is published by the American Association for the Advancement of Science, 1200 New York Avenue NW, Washington, DC 20005. 2017 © The Authors, some rights reserved; exclusive licensee American Association for the Advancement of Science. No claim to original U.S. Government Works. The title *Science Advances* is a registered trademark of AAAS.