Generative Quantum Machine Learning via Denoising Diffusion Probabilistic Models

Bingzhi Zhang,[1,2] Peng Xu,[3] Xiaohui Chen,[4,*] and Quntao Zhuang[2,1,†]
[1]Department of Physics and Astronomy, University of Southern California, Los Angeles, California 90089, USA
[2]Ming Hsieh Department of Electrical and Computer Engineering, University of Southern California, Los Angeles, California 90089, USA
[3]Department of Statistics, University of Illinois at Urbana-Champaign, Champaign, Illinois 61820, USA
[4]Department of Mathematics, University of Southern California, Los Angeles, California 90089, USA

(Received 18 November 2023; accepted 31 January 2024; published 5 March 2024)

Deep generative models are key-enabling technology to computer vision, text generation, and large language models. Denoising diffusion probabilistic models (DDPMs) have recently gained much attention due to their ability to generate diverse and high-quality samples in many computer vision tasks, as well as to incorporate flexible model architectures and a relatively simple training scheme. Quantum generative models, empowered by entanglement and superposition, have brought new insight to learning classical and quantum data. Inspired by the classical counterpart, we propose the *quantum denoising diffusion probabilistic model* (QuDDPM) to enable efficiently trainable generative learning of quantum data. QuDDPM adopts sufficient layers of circuits to guarantee expressivity, while it introduces multiple intermediate training tasks as interpolation between the target distribution and noise to avoid barren plateau and guarantee efficient training. We provide bounds on the learning error and demonstrate QuDDPM's capability in learning correlated quantum noise model, quantum many-body phases, and topological structure of quantum data. The results provide a paradigm for versatile and efficient quantum generative learning.

DOI: 10.1103/PhysRevLett.132.100602

Variational parametrized quantum circuits (PQCs) [1–4] provide a near-term platform for quantum machine learning [5–7]. In terms of generative models [8–11], quantum generative adversarial networks (QuGANs) have been recently proposed [12–16], in analogy to classical generative adversarial networks (GANs) [17]. Despite the success, classical GAN models are known for training issues such as mode collapse. In classical deep learning, denoising diffusion probabilistic models (DDPMs) and their close relatives [18–22] have recently gained much attention due to relatively simple training schemes and their ability to generate diverse and high-quality samples in many computer vision tasks [23–26] over the best GANs, and to incorporate flexible model architectures.

In this Letter, we propose the *quantum denoising diffusion probabilistic model* (QuDDPM) as an efficiently trainable scheme to generative quantum learning, through a coordination between a forward noisy diffusion process via quantum scrambling [27,28] and a backward denoising process via quantum measurement. We provide bounds on the learning error and then demonstrate QuDDPM's capability in examples relevant to characterizing quantum device noises, learning quantum many-body phases, and capturing topological structure of quantum data. For an n-qubit problem, QuDDPM adopts linear-in-n layers of circuits to guarantee expressivity, while it introduces $T \sim n/\log(n)$ intermediate training tasks to guarantee efficient training.

General formulation of QuDDPM.—We consider the task of the generating new elements from an unknown distribution $\mathcal{E}_0$ of quantum states, provided only a number of samples $\mathcal{S}_0 = \{|\psi_k\rangle\} \sim \mathcal{E}_0$ from the distribution. The task under consideration—generating individual states from the distribution (e.g., a single Haar random state or K-design state)—is *not* equivalent to generating the average state of a distribution (e.g., a fully mixed state for Haar ensemble) considered in previous works of QuGAN [12]. To complete the task, QuDDPM learns a map from a noisy unstructured distribution of states to the structured target distribution $\mathcal{E}_0$. It does so via a divide-and-conquer strategy of creating smooth interpolations between the target distribution and full noise, so that the training is divided to subtasks on a low-depth circuit to avoid barren plateau [29–32].

As shown in Fig. 1, QuDDPM includes two quantum circuits, one to enable the forward diffusion of sample data toward noise via scrambling and one to realize the backward denoising from noise toward generated data via measurement. For each data $|\psi_i^{(0)}\rangle$, the forward scrambling circuit [Fig. 1(a)] samples a series of T random unitary gates $U_1^{(i)}, \ldots, U_T^{(i)}$ independently, such that the ensemble $\mathcal{S}_k = \{|\psi_i^{(k)}\rangle = \prod_{\ell=1}^{k} U_\ell^{(i)} |\psi_i^{(0)}\rangle\}_i$ evolves from the sample data toward a random ensemble of pure states from $k = 0$ to $k = T$. A Bloch sphere visualization of such a forward scrambling dynamics is depicted in Figs. 1(b1)–1(b5) for a

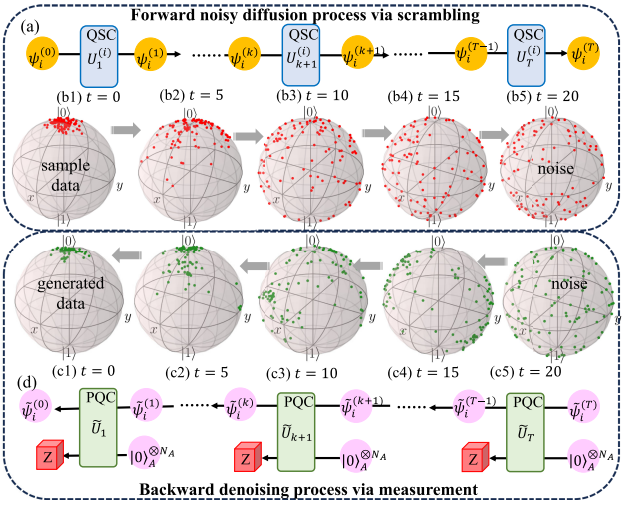

FIG. 1. Schematic of QuDDPM. The forward noisy process is implemented by a quantum scrambling circuit (QSC) in (a), while in the backward denoising process is achieved via measurement enabled by ancilla and PQC in (d). Subplots (b1)–(b5) and (c1)–(c5) present the Bloch sphere dynamics in generation of states clustering around $|0\rangle$, where convergence can be seen despite sample fluctuations, as shown in the Supplemental Material [33].

toy problem of learning single-qubit states $\mathcal{S}_0$ clustered around a single pure state, e.g., $|0\rangle$, (b1), where the noise $\mathcal{S}_T$ is uniform on the Bloch sphere (b5).

With the interpolation from data $\mathcal{S}_0$ and noise $\mathcal{S}_T$ in hand, the backward process can start from randomly sampled noise $\tilde{\mathcal{S}}_T$ [Fig. 1(c5)] and reduce the noise gradually via measurement step by step, toward the final generated data $\tilde{\mathcal{S}}_0$ [Fig. 1(c1)] that mimic the sample data [Fig. 1(b1)]. Measurements are necessary, as the denoising map is contractive and maintains the purity of all generated data in $\tilde{\mathcal{S}}_0$. As shown in Fig. 1(d), each denoising step adopts a unitary $\tilde{U}_k$ on the system plus n_A number of ancilla qubits in $|0\rangle$ and performs a projective measurement in computational bases on the ancilla after the unitary $\tilde{U}_k$. Starting from the state $|\tilde{\psi}_i^{(T)}\rangle$, which is randomly sampled from noise ensemble, each unitary plus measurement step evolves the random state toward the generated data $|\tilde{\psi}_i^{(0)}\rangle$. Note that here all unitaries $\tilde{U}_k$ are fixed after training. In practice, the generation of noisy $|\tilde{\psi}_i^{(T)}\rangle$ can be directly completed by running the T layers of the forward scrambling circuit on an arbitrary initial state. Via training, the denoising process learns information about the target from the ensembles in the forward scrambling, stores information in the circuit parameters, and then encodes onto the generated data.

Training strategy.—In classical DDPM, the Gaussian nature of the diffusion allows efficient training via maximizing an evidence lower bound for the log-likelihood function, which can be evaluated *analytically* [18,19,33]. However, in QuDDPM, we do not expect such analytical simplification to exist at all—classical simulation of quantum device is inherently inefficient. Instead, the training of the QuDDPM relies on the capability of quantum measurements to extract information about the ensemble of quantum states for the efficient evaluation of a loss function.

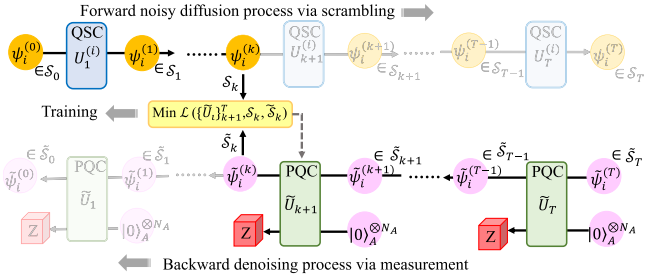

FIG. 2. The training of QuDDPM at each step $t = k$. Pairwise distance between states in generated ensemble $\tilde{\psi}_i^{(k)} \in \tilde{\mathcal{S}}_k$ and true diffusion ensemble $\psi_j^{(k)} \in \mathcal{S}_k$ is measured and utilized in the evaluation of the loss function $\mathcal{L}$.

The training of a T-step QuDDPM consists of T training cycles, starting from the first denoising step $\tilde{U}_T$ toward the last $\tilde{U}_1$. As shown in Fig. 2, at the training cycle $(T + 1 - k)$, the forward noisy diffusion process is implemented from $U_1^{(i)}$ to $U_k^{(i)}$ to generate the noisy ensemble $\mathcal{S}_k = \{|\psi_i^{(k)}\rangle\}_i$, while the backward denoising process performs the denoising steps $\tilde{U}_T$ to $\tilde{U}_{k+1}$ to generate the denoising ensemble $\tilde{\mathcal{S}}_k = \{|\tilde{\psi}_i^{(k)}\rangle\}_i$. Within the training cycle, the parameters of the denoising PQC $\tilde{U}_{k+1}$ are updated such that the generated denoising ensemble $\tilde{\mathcal{S}}_k$ converges to the noisy ensemble $\mathcal{S}_k$. Therefore, QuDDPM divides the original training problem into T smaller and easier ones. Indeed, even with a global loss function, for n qubits we can divide the $\Omega(n)$ layers (required by expressivity) of gates into $T \in \Omega(n/\log n)$ diffusion steps, such that each $\tilde{U}_{k+1}$ has order $\log(n)$ layers of gates to avoid barren plateau [30].

Loss function.—To enable training, a loss function quantifies the distance between the two ensembles of quantum states. In this Letter, we focus on the maximum mean discrepancy (MMD) [46] and the Wasserstein distance [47,48] based on the state overlaps $|\langle\psi_i^{(k)}|\tilde{\psi}_i^{(k)}\rangle|^2$ estimated via a swap test [33].

Given two independent distributions of pure states $\mathcal{E}_1$ and $\mathcal{E}_2$ on the state vector space V, the *statewise fidelity* between $|\psi\rangle$ and $|\phi\rangle$ is defined as $F(|\phi\rangle, |\psi\rangle) = |\langle\phi|\psi\rangle|^2$, and we can further define the *mean fidelity* as

$$\bar{F}(\mathcal{E}_1, \mathcal{E}_2) = \mathbb{E}_{|\phi\rangle\sim\mathcal{E}_1, |\psi\rangle\sim\mathcal{E}_2}[|\langle\phi|\psi\rangle|^2], \tag{1}$$

where the random states $|\phi\rangle \sim \mathcal{E}_1$ and $|\psi\rangle \sim \mathcal{E}_2$ are drawn independently. Since the fidelity F is a symmetric and positive definite quadratic kernel, according to the theory of reproducing kernel Hilbert space [49], the MMD distance can be written as

$$\mathcal{D}_{\mathrm{MMD}}(\mathcal{E}_1,\mathcal{E}_2)=\bar{F}(\mathcal{E}_1,\mathcal{E}_1)+\bar{F}(\mathcal{E}_2,\mathcal{E}_2)-2\bar{F}(\mathcal{E}_2,\mathcal{E}_1), \quad (2)$$

which allows the estimation of MMD through sampled state ensembles $\mathcal{S}_1$ and $\mathcal{S}_2$. The expressivity of general MMD as a statistical distance measure depends on the kernel. On one hand, identifiability requires that the distance be zero if and only if $\mathcal{E}_1=\mathcal{E}_2$. On the other hand, one also needs to ensure the quality of statistical estimation of the distance with a finite sample size of state ensembles. Hence, whether fidelity (1) is a proper kernel choice is problem-dependent. In the Supplemental Material [33], we show an example where $\mathcal{D}_{\mathrm{MMD}}$ in Eq. (2) fails to distinguish two simple distributions. To resolve such an issue, we may alternatively consider the Wasserstein distance, a geometrically meaningful distance for comparing complex data distributions based on the theory of optimal transport [47,48] (see the Appendix for details).

As shown in Fig. 2, in the training cycle $t=k$, loss is a function of the unitary $\{\tilde{U}_\ell\}_{k+1}^{T}$ and depends on the noise distribution $\tilde{\mathcal{E}}_T$ and the scrambled data distribution $\mathcal{E}_k$,

$$\mathcal{L}(\{\tilde{U}_\ell\}_{k+1}^{T},\mathcal{E}_k,\tilde{\mathcal{E}}_T)=\mathcal{D}(\mathcal{E}_k,\tilde{\mathcal{E}}_k[\{\tilde{U}_\ell\}_{k+1}^{T},\tilde{\mathcal{E}}_T]), \quad (3)$$

where $\mathcal{D}$ can be the MMD distance or the Wasserstein distance. The distribution $\tilde{\mathcal{E}}_k$ is a function of all the reverse denoising steps from T to $k+1$ and the noise distribution $\tilde{\mathcal{E}}_T$. In practice, we use finite samples to approximate the loss function as $\mathcal{L}(\{\tilde{U}_\ell\}_{k+1}^{T},\mathcal{S}_k,\tilde{\mathcal{S}}_T)$.

The toy example in Fig. 1 adopted the MMD distance in the loss function, and details of the training can be found in the Supplemental Material [33] (see Ref. [50] for codes and data, and Table I for details of parameters). Here we present the training history of a more challenging 2-qubit example of preparing states clustered around $|0,0\rangle$, to allow a meaningful comparison with other algorithms. In each of the 20 steps of training cycles, the loss function is minimized till convergence. To quantify the convergence, we also evaluate the MMD distance (see Fig. 3) $\mathcal{D}(\tilde{\mathcal{S}}_t,\mathcal{E}_0)$ between the true distribution $\mathcal{E}_0$ and the trained ensemble of states $\tilde{\mathcal{S}}_t$ throughout the training cycles (blue), showing a convergence toward $\mathcal{D}=0$. The periodic spikes show the initial increase of the MMD distance at each training cycle, due to introducing a randomly initialized PQC in a new denoising step. For reference, we also plot the evolution of the MMD distance throughout the forward-diffusion (red circles), which starts from zero at diffusion step 0 and grows toward a larger value as the diffusion step increases (from right to left). We see the training results (blue) follow closely to the diffusion results (red) as expected. In addition, the testing results (green) also agree well with the training results (blue) for QuDDPM.

As benchmarks, we consider two major quantum generative models, QuGAN and quantum direct transport (QuDT). QuDT can be regarded as the generalization of the quantum circuit Born machine [51–54] toward quantum data. Previous works of both models only considered a single quantum state or classical distributions [14,52,55], and here we generalize them to adapt to the state ensemble generation task by allowing Haar random states as inputs and introducing ancilla to be measured [33]. For a fair comparison, we keep the number of variational parameters of generator circuits in QuDT and QuGAN the same as QuDDPM, listed in Table II. As shown in Fig. 3, QuDT and QuGAN converge to generate ensembles with a substantial MMD deviation to the true ensemble, demonstrating QuDDPM's advantage due to its unique diffusion and denoising process.

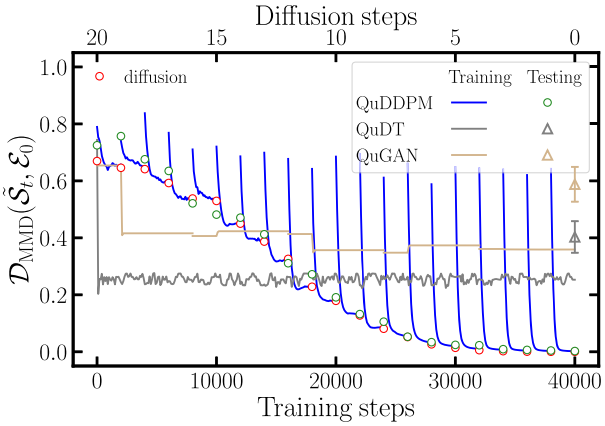

FIG. 3. The decay of MMD distance $\mathcal{D}$ between generated ensemble $\tilde{\mathcal{S}}_t$ using different models and target ensemble of states $\mathcal{E}_0$ clustered around $|0,0\rangle$ versus training steps. The converged value is $\mathcal{D}\simeq 0.002$ for QuDDPM, showing an advantage of 2 orders of magnitude over QuDT and QuGAN.

Gate complexity and convergence.—Now we discuss the number of local gates required and convergence analysis for QuDDPM to solve an n-qubit generative task. For simplicity, we assume the qubits are one-dimensional with nearest-neighbor interactions, while similar counting can be done for other cases. To guarantee convergence toward noise, the forward scrambling circuits need a linear number of layers in n as predicted by K design [56,57], leading to $O(n^2)$ total gates. The backward circuit will be similar, with at most $n_A\le 2n$ additional ancillas and $O(n^2)$ gates, leading the overall gate complexity of QuDDPM to be $O(n^2)$.

Similar to the classical case [58], the total error of QuDDPM involves three parts,

$$E\simeq E_{\mathrm{diff}}+E_{\mathcal{M}}+E_{\mathrm{gen}}, \quad (4)$$

with a deviation E_{diff} of $\mathcal{S}_T$ to true random states, measurement error $E_{\mathcal{M}}$, and generalization error E_{gen}. We discuss the scaling of three parts separately in the following. Suppose the diffusion circuits approach an approximate K design; its diffusion error is known as [56]

$$E_{\mathrm{diff}} \sim 2^{nK} e^{-T/A(K)C} \sim \mathcal{O}(e^{-T}), \tag{5}$$

where $A(K) = \lceil \log_2(4K) \rceil^2 K^5 t^{3.1/\log(2)}$ is a polynomial of K and C is a constant determined by the circuit in a single step. For measurement, the standard error in estimating the fidelity F_{ij} between any two states $|\psi_i\rangle, |\tilde{\psi}_j\rangle$ is $\mathrm{SE}(F_{ij}) = \sqrt{(1-F_{ij})/m}$, where m is the number of repetitions of measurement. With N data in the two sets $\mathcal{S}, \tilde{\mathcal{S}}$, the measurement error of estimating the *mean fidelity* is

$$E_{\mathcal{M}} = \frac{1}{N^2}\sqrt{\sum_{i,j=1}^{N} \mathrm{SE}(F_{ij})^2} \sim \mathcal{O}\left(\frac{1}{N\sqrt{m}}\right). \tag{6}$$

Finally, we provide numerical evidence that the generalization error [59,60]

$$E_{\mathrm{gen}}(\{\tilde{U}_\ell\}_1^T) \equiv \mathcal{L}(\{\tilde{U}_\ell\}_1^T, \mathcal{E}_0, \tilde{\mathcal{E}}_T) - \mathcal{L}(\{\tilde{U}_\ell\}_1^T, \mathcal{S}_0, \tilde{\mathcal{S}}_T) \tag{7}$$

has the scaling $\mathcal{O}(1/TN)$, as shown in Fig. 4 for an $n = 4$ qubit clustering state generation task. Here we estimate the generalization error via a validation set independently sampled, while the proof is an open problem [59,60]. The $1/N$ scaling agrees with classical machine learning results [61,62].

Applications.—To showcase QuDDPM's applications, we consider a particular realization of QuDDPM with each unitary $U_k^{(i)}$ and $\tilde{U}_k$ implemented by the fast scrambling model [63]—layers of general single-qubit rotations in between homogeneous tunable entangling layers of all-to-all ZZ rotations—and hardware efficient ansatz [64]—layers of X and Y single-qubit rotations in between layers of nearest-neighbor control-Z gates—separately [33]. While the MMD distance characterization similar to Fig. 3 is presented in the Supplemental Material [33], we adopt more direct measures of performance in each application.

Learning correlated noise.—When a real quantum device is programmed to generate a quantum state, it inevitably suffers from potentially correlated errors in the gate control parameters [65,66]. As a result, the generated states $\mathcal{S}_0$ are close to the target state but have nontrivial coherent errors, which can be learned by QuDDPM. We take a 2-qubit example of the target state $|\Psi\rangle = c_0|00\rangle + c_1|01\rangle + c_3|11\rangle$ under the influence of fully correlated noise, where $e^{-i\delta X_1 X_2}$ and $e^{-i\delta Z_1 Z_2}$ rotations happen with probability p and $1-p$. Here X_k and Z_k are Pauli operators for qubit k. In each case, the angle of rotation δ is uniformly sampled from the range $[-\delta_0, \delta_0]$. As the $|10\rangle$ component in the superposition only appears when XX error happens, we can utilize average fidelity $\overline{F_{10}} = \mathbb{E}_{\tilde{\mathcal{S}}_0}|\langle 10|\psi^{(0)}\rangle|^2$ as the performance metric to estimate the error probability p via $\tilde{p} = \overline{F_{10}}/(|c_1|^2 \mathbb{E}_\delta[\sin^2\delta])$. We show a numerical example in Fig. 5(a), where the generated ensemble average fidelity in training and testing agrees with the theoretical prediction up to a finite sample size deviation.

Learning many-body phases.—As a proof of principle, we take the simple and well-known transverse-field Ising model (TFIM) described by the Hamitonian $H_{\mathrm{TFIM}} = -\sum_i Z_i Z_{i+1} - g\sum_i X_i$. When g increases from zero, the system undergoes a phase transition from an ordered ferromagnetic phase to a disordered phase, with the critical point at $g = 1$. The states before diffusion are chosen from ground states of H_{TFIM} with $g \in [0.2, 0.4)$ uniformly distributed. To test the capability of QuDDPM, we utilize the magnetization, $M = (\sum_i Z_i)/n$, to identify the phase of generated states from QuDDPM, and show its distribution in Fig. 5(b). Most generated states (blue and green) of QuDDPM live in the ferromagentic phase, and show a sharp contrast to the random states (orange).

Learning nontrivial topology.—We consider the ensemble of states with a ring structure—generated by applying a unitary on a single state, e.g., $|\psi_i\rangle = e^{-i\boldsymbol{x}_i\cdot\boldsymbol{G}}|0\rangle$, which models the scenario where one encodes the classical data $\boldsymbol{x}_i$ onto the quantum data ψ_i, as commonly adopted in

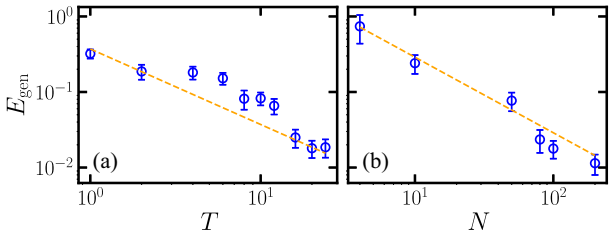

FIG. 4. The generalization error of QuDDPM in generating cluster states versus (a) diffusion steps T and (b) training dataset size N. Dots are numerical results, and orange dashed line is linear fitting results with both exponents equal to 1 within the numerical precision.

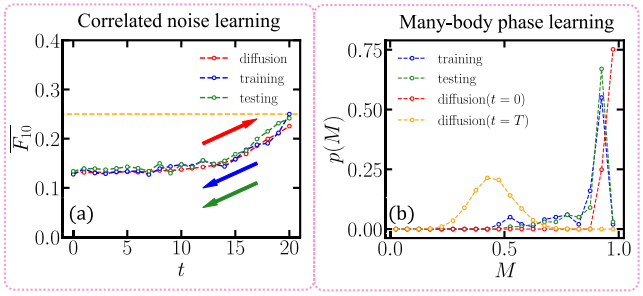

FIG. 5. Generation of states with probabilistic correlated noise on a specific state in (a) and (b) states with ferromagnetic phase. In (a), average fidelity $\overline{F_{10}}$ between states at step t and $|10\rangle$ for diffusion (red), training (blue), and testing (green) are plotted. In (b), we show the distribution of magnetization for generated data from training (blue) and testing (green) dataset, and compared to true data (red) and full noise (orange). Four qubits are considered in (b).

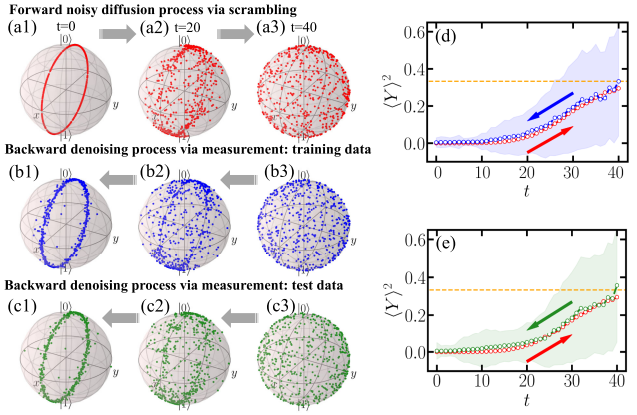

FIG. 6. Bloch visualization of the forward (a1)–(a3) and backward (b1)–(b3),(c1)–(c3) process. (d),(e) deviation of generated states from unit circle in X-Z plane. The deviation $\langle Y\rangle^2$ for forward diffusion (red), backward training (blue), and backward test (green) are plotted. The shaded area shows the sample standard deviation.

quantum machine learning to solve classical problems [15,67–69]. We test QuDDPM with a single qubit toy example, where the generators are chosen as Pauli Y and the rotation angles are uniform in $[0, 2\pi)$. In the QuDDPM training, we use the Wasserstein distance [48] to cope with the nontrivial topology. The forward noisy diffusion process on the sample data and the backward denoising process for training and testing are depicted in Fig. 6. To quantitatively evaluate the performance of QuDDPM, we evaluate the deviation by Pauli Y expectation $\langle Y\rangle^2$ in Figs. 6(d) and 6(e), where gradual transition between zero and a Haar value of $1/3$ is observed in both forward diffusion and backward denoising.

Q. Z. and B. Z. acknowledge support from NSF CAREER Awards No. CCF-2240641 and No. ONR N00014-23-1-2296. X. C. acknowledges support from NSF CAREER Award No. DMS-2347760.

Appendix A: On details of parameters.—We list hyperparameters and performance for all generative learning tasks in Table I for reference, and state the targeted distribution of states to generate in the following. The major codes and data of the work can be found in Ref. [50].

In Figs. 1(b) and 1(c), we consider data in the form of $|\psi^{(0)}\rangle \sim |0\rangle + \epsilon c_1|1\rangle$ up to a normalization constant where $\Re c_1, \Im c_1 \sim \mathcal{N}(0, 1)$ is Gaussian distributed, and the scale factor is chosen as $\epsilon = 0.08$. We have taken single qubit rotations as $U_k^{(i)}$, where each angle is randomly sampled, e.g., from the uniform distribution $U[-\pi/8, \pi/8]$. In the generation of states with probabilistic correlated noise of Fig. 5(a), the noise perturbation range is $\delta \in [-\pi/3, \pi/3]$.

TABLE I. List of hyperparameters of quantum denoising diffusion probabilistic model (QuDDPM) and its performance in different generative learning tasks. To test the performance after training, we randomly sample N_{te} random noise states, and perform the optimized backward PQC to generate the sampled data. Dataset size $N_{\text{tr}} = N_{\text{te}} = N$. n is the number of data qubit n and n_A is the ancilla qubit. L is the PQC depth. T is the diffusion steps. For cluster state generation, we evaluate the average fidelity with the center state in each cluster, i.e., $|0\rangle$ for single qubit and $|0, 0\rangle$ for 2 qubit [33]. We also specify the cost function, among the two choices: maximum 131 mean discrepancy (MMD) and Wasserstein distance.

Generation task	n	n_A	L	T	N	Cost function	Performance
Clustered state (Fig. 1 of main text and Figs. 4 and 6a of Ref. [33])	1	1	4	20	100	MMD	$\overline{F_{0,\text{data}}} = 0.987 \pm 0.013$ $\overline{F_{0,\text{tr}}} = 0.992 \pm 0.021$ $\overline{F_{0,\text{te}}} = 0.993 \pm 0.014$
Clustered state (Fig. 3 of main text and Fig. 6(b) of Ref. [33])	2	1	6	20	100	MMD	$\overline{F_{0,\text{data}}} = 0.977 \pm 0.014$ $\overline{F_{0,\text{tr}}} = 0.952 \pm 0.070$ $\overline{F_{0,\text{te}}} = 0.944 \pm 0.075$
Clustered state [Figs. 4(a) and 4(b)]	4	2	8	20 [Fig. 4(b)]	100 [Fig. 4(a)]	MMD	See Figs. 4(a) and 4(b)
Correlated noise [Fig. 5(a)]	2	2	6	20	500	MMD	Data: 0.129 Training: 0.128 Testing: 0.133
Many-body phase [Fig. 5(b)]	4	2	12	30	100	MMD	Measured by magnetization. Data: 1 Training: 0.9 Testing: 0.96
Circular states (Fig. 6)	1	2	6	40	500	Wasserstein	Data: $\overline{\langle Y\rangle^2} = 0$ Training: $\overline{\langle Y\rangle^2} = 0.00367 \pm 0.0251$ Testing: $\overline{\langle Y\rangle^2} = 0.00506 \pm 0.0439$

TABLE II. List of hyperparameters of quantum denoising diffusion probabilistic model (QuDDPM), quantum direct transport (QuDT), and quantum generative adversarial networks (QuGANs) for generating a clustered state in Figs. 3 and 8 in Ref. [33]. Dataset size $N_{\mathrm{tr}} = N_{\mathrm{te}} = N$. n is the number of data qubit n, and n_A is the ancilla qubit. In performance, the mean fidelity with the center state of the cluster $|0,0\rangle$ is $\overline{F_0} = \mathbb{E}_{|\psi\rangle \in \tilde{S}} |\langle 0,0|\psi\rangle|^2$, and for true data it is $\overline{F_{0,\mathrm{data}}} = 0.977 \pm 0.014$ [33].

Model	n	n_A	No. variational parameters	N	Cost function	Performance
QuDDPM	2	1	720	100	MMD	$\overline{F_{0,\mathrm{tr}}} = 0.947 \pm 0.070$, $\overline{F_{0,\mathrm{te}}} = 0.948 \pm 0.061$
QuDT	2	1	720	100	MMD	$\overline{F_{0,\mathrm{tr}}} = 0.572 \pm 0.321$ $\overline{F_{0,\mathrm{te}}} = 0.465 \pm 0.349$
QuGAN	2	1	720 (generator) 96 (discriminator)	100	Error probability based cost function	$\overline{F_{0,\mathrm{tr}}} = 0.570 \pm 0.250$ $\overline{F_{0,\mathrm{te}}} = 0.443 \pm 0.269$

Appendix B: On Wasserstein distance.—For pure states, choosing the quantum trace distance (equaling infidelity) $D^2(|\phi\rangle, |\psi\rangle) = 1 - |\langle\phi|\psi\rangle|^2$, then Kantorovich's formulation for optimal transportation involves solving the following optimization problem:

$$\mathsf{OPT} := \min_{\pi \in \Pi(\mathcal{E}_1, \mathcal{E}_2)} \int_{V \times V} D^p(|\phi\rangle, |\psi\rangle) d\pi(|\phi\rangle, |\psi\rangle) \tag{B1}$$

for $p \geq 1$, where $\Pi(\mathcal{E}_1, \mathcal{E}_2)$ is the set of admissible *transport plans* (i.e., *couplings*) of probability distributions on $V \times V$ such that $\pi(B \times V) = \mathcal{E}_1(B)$ and $\pi(V \times B) = \mathcal{E}_2(B)$ for any measurable $B \subset V$; namely $\Pi(\mathcal{E}_1, \mathcal{E}_2)$ stands for all distributions with marginals as $\mathcal{E}_1$ and $\mathcal{E}_2$. The Kantorovich problem in Eq. (B1) induces a metric, known as the *p-Wasserstein distance*, on the space $\mathcal{P}_p(V)$ of probability distributions on V with finite pth moment. In particular, the p-Wasserstein distance $W_p(\mathcal{E}_1, \mathcal{E}_2) = \mathsf{OPT}^{1/p}$, and it has identifiability in the sense that $W_p(\mathcal{E}_1, \mathcal{E}_2) = 0$ if and only if $\mathcal{E}_1 = \mathcal{E}_2$. More details can be found in Ref. [33].

Appendix C: On related works.—The proposed QuDDPM represents an application of the theoretical idea of quantum information scrambling [27,28] in the forward diffusion, and its backward denoising also connects to the measurement-induced phase transitions [70]. Here we point out that our forward diffusion circuits include an actual implementation of scrambling as part of the QuDDPM algorithm, while previous papers utilize tools from the study of quantum scrambling to understand quantum neural networks [71,72].

Below, we discuss several related works. Reference [73] utilizes a diffusion map (DM) for unsupervised learning of topological phases, and [74] proposes a diffusion K-means manifold clustering approach based on the diffusion distance [75]. A quantum DM algorithm has also been considered [76] for potentially quantum speedup. However, these works are not on generative learning and do not consider any denoising process. Layerwise training [77] also attempts to divide a training problem into subtasks in nongenerative learning; however, the performance of such strategies is limited [78]. QuDDPM integrates the division of training task and an actual noisy diffusion process to enable provable benefits in training. After the completion of our Letter, we became aware of a recent paper [79], where hybrid quantum-classical DDPM is proposed. Our Letter focuses on quantum data; provides explicit construction of quantum diffusion and quantum denoising, loss function, training strategy and error analyses; and presents several applications.

Appendix D: On future directions.—Finally, we point out some future directions, besides various applications of QuDDPM in learning quantum systems. Our current QuDDPM architecture requires a loss function based on fidelity estimations. For large systems, fidelity estimation can be challenging to implement. Toward efficient training in large systems, alternative loss functions can be adopted. For example, one may consider adopting another quantum circuit trained for telling the ensembles apart, such as a quantum convolutional neural network [4] and other circuit architecture [80]. Such an approach will combine QuDDPM and the adversarial agent in QuGAN to resolve the training problem in QuGAN. Another future direction is controlled diffusion [81]: when the ensemble has special symmetry, one can restrict the forward scrambling, the backward denoising, and the random noise ensemble to that symmetry. It is also an interesting open problem of how to introduce a control knob such that QuDDPM can learn multiple distributions and generate states according to an input requesting one of the distributions.

Besides learning quantum errors and many-body phases, quantum sensor networks [82,83] provide another application scenario of QuDDPM. In this scenario, one sends quantum probes to sense a unitary physical process; on the return side, the receiver will collect a pure state from a distribution in the ideal case. It is an open problem of how

QuDDPM can be adopted to provide an advantage in quantum sensing.

Furthermore, the data can also be quantum states encoding classical data, where QuDDPM can also process classical data. Benchmarking QuDDPM and previous algorithms for classical data generative learning is an open direction.

*xiaohuic@usc.edu
†qzhuang@usc.edu

[1] M. Cerezo, A. Arrasmith, R. Babbush, S. C. Benjamin, S. Endo, K. Fujii, J. R. McClean, K. Mitarai, X. Yuan, L. Cincio *et al.*, Nat. Rev. Phys. **3**, 625 (2021).
[2] N. Killoran, T. R. Bromley, J. M. Arrazola, M. Schuld, N. Quesada, and S. Lloyd, Phys. Rev. Res. **1**, 033063 (2019).
[3] A. Abbas, D. Sutter, C. Zoufal, A. Lucchi, A. Figalli, and S. Woerner, Nat. Comput. Sci. **1**, 403 (2021).
[4] I. Cong, S. Choi, and M. D. Lukin, Nat. Phys. **15**, 1273 (2019).
[5] J. Biamonte, P. Wittek, N. Pancotti, P. Rebentrost, N. Wiebe, and S. Lloyd, Nature (London) **549**, 195 (2017).
[6] P. Rebentrost, M. Mohseni, and S. Lloyd, Phys. Rev. Lett. **113**, 130503 (2014).
[7] S. Lloyd, M. Mohseni, and P. Rebentrost, Nat. Phys. **10**, 631 (2014).
[8] X. Gao, E. R. Anschuetz, S.-T. Wang, J. I. Cirac, and M. D. Lukin, Phys. Rev. X **12**, 021037 (2022).
[9] A. Khoshaman, W. Vinci, B. Denis, E. Andriyash, H. Sadeghi, and M. H. Amin, Quantum Sci. Technol. **4**, 014001 (2018).
[10] M. H. Amin, E. Andriyash, J. Rolfe, B. Kulchytskyy, and R. Melko, Phys. Rev. X **8**, 021050 (2018).
[11] X. Gao, Z.-Y. Zhang, and L.-M. Duan, Sci. Adv. **4**, eaat9004 (2018).
[12] S. Lloyd and C. Weedbrook, Phys. Rev. Lett. **121**, 040502 (2018).
[13] P.-L. Dallaire-Demers and N. Killoran, Phys. Rev. A **98**, 012324 (2018).
[14] L. Hu, S.-H. Wu, W. Cai, Y. Ma, X. Mu, Y. Xu, H. Wang, Y. Song, D.-L. Deng, C.-L. Zou *et al.*, Sci. Adv. **5**, eaav2761 (2019).
[15] H.-L. Huang, Y. Du, M. Gong, Y. Zhao, Y. Wu, C. Wang, S. Li, F. Liang, J. Lin, Y. Xu *et al.*, Phys. Rev. Appl. **16**, 024051 (2021).
[16] E. Y. Zhu, S. Johri, D. Bacon, M. Esencan, J. Kim, M. Muir, N. Murgai, J. Nguyen, N. Pisenti, A. Schouela *et al.*, Phys. Rev. Res. **4**, 043092 (2022).
[17] I. Goodfellow *et al.*, Adv. Neural Inf. Process. Syst. **27** (2014), https://papers.nips.cc/paper_files/paper/2014/hash/5ca3e9b122f61f8f06494c97b1afccf3-Abstract.html.
[18] J. Sohl-Dickstein, E. Weiss, N. Maheswaranathan, and S. Ganguli, in *Proceedings of the 32nd International Conference on Machine Learning* (PMLR, 2015), pp. 2256–2265, https://proceedings.mlr.press/v37/sohl-dickstein15.html.
[19] J. Ho, A. Jain, and P. Abbeel, Adv. Neural Inf. Process. Syst. **33**, 6840 (2020), https://proceedings.neurips.cc/paper/2020/hash/4c5bcfec8584af0d967f1ab10179ca4b-Abstract.html.
[20] Y. Song, J. Sohl-Dickstein, D. P. Kingma, A. Kumar, S. Ermon, and B. Poole, in *Proceedings of the The Ninth International Conference on Learning Representations* (2021).
[21] A. Schneuing, Y. Du, C. Harris, A. Jamasb, I. Igashov, W. Du, T. Blundell, P. Lió, C. Gomes, M. Welling, M. Bronstein, and B. Correia, arXiv:2210.13695.
[22] Y. Chen, T. T. Georgiou, and M. Pavon, J. Optim. Theory Appl. **169**, 671 (2016).
[23] P. Dhariwal and A. Nichol, Adv. Neural Inf. Process. Syst. **34**, 8780 (2021), https://papers.nips.cc/paper/2021/hash/49ad23d1ec9fa4bd8d77d02681df5cfa-Abstract.html.
[24] G. Müller-Franzes, J. M. Niehues, F. Khader, S. T. Arasteh, C. Haarburger, C. Kuhl, T. Wang, T. Han, S. Nebelung, J. N. Kather *et al.*, Sci. Rep. **13**, 12098 (2023).
[25] A. Jalal, M. Arvinte, G. Daras, E. Price, A. G. Dimakis, and J. Tamir, Adv. Neural Inf. Process. Syst. **34**, 14938 (2021), https://proceedings.neurips.cc/paper_files/paper/2021/hash/7d6044e95a16761171b130dcb476a43e-Abstract.html.
[26] Y. Song and S. Ermon, Adv. Neural Inf. Process. Syst. **32** (2019), https://proceedings.neurips.cc/paper_files/paper/2019/hash/3001ef257407d5a371a96dcd947c7d93-Abstract.html.
[27] A. Nahum, J. Ruhman, S. Vijay, and J. Haah, Phys. Rev. X **7**, 031016 (2017).
[28] A. Nahum, S. Vijay, and J. Haah, Phys. Rev. X **8**, 021014 (2018).
[29] J. R. McClean, S. Boixo, V. N. Smelyanskiy, R. Babbush, and H. Neven, Nat. Commun. **9**, 4812 (2018).
[30] M. Cerezo, A. Sone, T. Volkoff, L. Cincio, and P. J. Coles, Nat. Commun. **12**, 1791 (2021).
[31] S. Wang, E. Fontana, M. Cerezo, K. Sharma, A. Sone, L. Cincio, and P. J. Coles, Nat. Commun. **12**, 6961 (2021).
[32] C. Ortiz Marrero, M. Kieferová, and N. Wiebe, PRX Quantum **2**, 040316 (2021).
[33] See Supplemental Material at http://link.aps.org/supplemental/10.1103/PhysRevLett.132.100602 for details on analyses, which include Refs. [34–45].
[34] B. D. Anderson, Stoch. Proc. Appl. **12**, 313 (1982).
[35] T. Hsing and R. Eubank, *Theoretical Foundations of Functional Data Analysis, with an Introduction to Linear Operators*, Wiley Series in Probability and Statistics (Wiley, New York, 2015).
[36] G. Peyré and M. Cuturi, Found. Trends Mach. Learn. **11**, 355 (2019).
[37] S. Chakrabarti, H. Yiming, T. Li, S. Feizi, and X. Wu, Adv. Neural Inf. Process. Syst. **32** (2019), https://proceedings.neurips.cc/paper_files/paper/2019/hash/f35fd567065af297ae65b621e0a21ae9-Abstract.html.
[38] G. De Palma, M. Marvian, D. Trevisan, and S. Lloyd, IEEE Trans. Inf. Theory **67**, 6627 (2021).
[39] B. T. Kiani, G. De Palma, M. Marvian, Z.-W. Liu, and S. Lloyd, Quantum Sci. Technol. **7**, 045002 (2022).
[40] J. A. Smolin and D. P. DiVincenzo, Phys. Rev. A **53**, 2855 (1996).
[41] S.-X. Zhang, J. Allcock, Z.-Q. Wan, S. Liu, J. Sun, H. Yu, X.-H. Yang, J. Qiu, Z. Ye, Y.-Q. Chen *et al.*, Quantum **7**, 912 (2023).
[42] J. R. Johansson, P. D. Nation, and F. Nori, Comput. Phys. Commun. **183**, 1760 (2012).

[43] R. Flamary *et al.*, J. Mach. Learn. Res. **22**, 1 (2021), https://jmlr.org/papers/v22/20-451.html.

[44] K. Huang, Z.-A. Wang, C. Song, K. Xu, H. Li, Z. Wang, Q. Guo, Z. Song, Z.-B. Liu, D. Zheng *et al.*, npj Quantum Inf. **7**, 165 (2021).

[45] M. Y. Niu, A. Zlokapa, M. Broughton, S. Boixo, M. Mohseni, V. Smelyanskyi, and H. Neven, Phys. Rev. Lett. **128**, 220505 (2022).

[46] A. Gretton, K. M. Borgwardt, M. J. Rasch, B. Schölkopf, and A. Smola, J. Mach. Learn. Res. **13**, 723 (2012), https://jmlr.csail.mit.edu/papers/v13/gretton12a.html.

[47] C. Villani, *Topics in Optimal Transportation*, Graduate studies in mathematics (American Mathematical Society, Providence, 2003).

[48] O. Oreshkov and J. Calsamiglia, Phys. Rev. A **79**, 032336 (2009).

[49] B. Schölkopfand A. J. Smola, *Learning with Kernels: Support Vector Machines, Regularization, Optimization, and Beyond* (MIT Press, 2018), 10.7551/mitpress/4175.001.0001.

[50] QuantGenMdl, howpublished = https://github.com/francis-hsu/quantgenmdl, note=Accessed: 2024-02-01.

[51] J.-G. Liu and L. Wang, Phys. Rev. A **98**, 062324 (2018).

[52] M. Benedetti, D. Garcia-Pintos, O. Perdomo, V. Leyton-Ortega, Y. Nam, and A. Perdomo-Ortiz, npj Quantum Inf. **5**, 45 (2019).

[53] B. Coyle, D. Mills, V. Danos, and E. Kashefi, npj Quantum Inf. **6**, 60 (2020).

[54] K. Gili, M. Hibat-Allah, M. Mauri, C. Ballance, and A. Perdomo-Ortiz, Quantum Sci. Technol. **8**, 035021 (2023).

[55] C. Zoufal, A. Lucchi, and S. Woerner, npj Quantum Inf. **5**, 103 (2019).

[56] F. G. Brandao, A. W. Harrow, and M. Horodecki, Commun. Math. Phys. **346**, 397 (2016).

[57] A. W. Harrow and S. Mehraban, Commun. Math. Phys. **401** (2023).

[58] S. Chen, S. Chewi, J. Li, Y. Li, A. Salim, and A. R. Zhang, arXiv:2209.11215.

[59] M. C. Caro, H.-Y. Huang, N. Ezzell, J. Gibbs, A. T. Sornborger, L. Cincio, P. J. Coles, and Z. Holmes, Nat. Commun. **14**, 3751 (2023).

[60] L. Banchi, J. Pereira, and S. Pirandola, PRX Quantum **2**, 040321 (2021).

[61] N. Srebro, K. Sridharan, and A. Tewari, Adv. Neural Inf. Process. Syst. **23** (2010), https://proceedings.neurips.cc/paper_files/paper/2010/hash/76cf99d3614e23eabab16fb27-e944bf9-Abstract.html.

[62] R. Yao, X. Chen, and Y. Yang, *Proceedings of Thirty Fifth Conference on Learning Theory* (PMLR, 2022), pp. 2242–2275, https://proceedings.mlr.press/v178/yao22a.html.

[63] R. Belyansky, P. Bienias, Y. A. Kharkov, A. V. Gorshkov, and B. Swingle, Phys. Rev. Lett. **125**, 130601 (2020).

[64] A. Kandala, A. Mezzacapo, K. Temme, M. Takita, M. Brink, J. M. Chow, and J. M. Gambetta, Nature (London) **549**, 242 (2017).

[65] R. Harper, S. T. Flammia, and J. J. Wallman, Nat. Phys. **16**, 1184 (2020).

[66] S. Chen, Y. Liu, M. Otten, A. Seif, B. Fefferman, and L. Jiang, Nat. Commun. **14**, 52 (2023).

[67] V. Havlíček, A. D. Córcoles, K. Temme, A. W. Harrow, A. Kandala, J. M. Chow, and J. M. Gambetta, Nature (London) **567**, 209 (2019).

[68] M. Schuld, arXiv:2101.11020.

[69] G. Li, R. Ye, X. Zhao, and X. Wang, Adv. Neural Inf. Process. Syst. **35**, 19456 (2022), https://proceedings.neurips.cc/paper_files/paper/2022/hash/7b2d0730df1edd8c97df4bf83696025d-Abstract-Conference.html.

[70] B. Skinner, J. Ruhman, and A. Nahum, Phys. Rev. X **9**, 031009 (2019).

[71] H. Shen, P. Zhang, Y.-Z. You, and H. Zhai, Phys. Rev. Lett. **124**, 200504 (2020).

[72] R. J. Garcia, K. Bu, and A. Jaffe, J. High Energy Phys. 03 (2022) 001.

[73] J. F. Rodriguez-Nieva and M. S. Scheurer, Nat. Phys. **15**, 790 (2019).

[74] X. Chen and Y. Yang, Appl. Comput. Harmon. Anal. **52**, 303 (2021).

[75] R. R. Coifman and S. Lafon, Appl. Comput. Harmon. Anal. **21**, 5 (2006), Special Issue: Diffusion Maps and Wavelets.

[76] A. Sornsaeng, N. Dangniam, P. Palittapongarnpim, and T. Chotibut, Phys. Rev. A **104**, 052410 (2021).

[77] A. Skolik, J. R. McClean, M. Mohseni, P. van der Smagt, and M. Leib, Quantum. Mach. **3**, 5 (2021).

[78] E. Campos, D. Rabinovich, V. Akshay, and J. Biamonte, Phys. Rev. A **104**, L030401 (2021).

[79] M. Parigi, S. Martina, and F. Caruso, arXiv:2308.12013.

[80] B. Zhang and Q. Zhuang, Quantum Sci. Technol. **7**, 035017 (2022).

[81] Y. Song, L. Shen, L. Xing, and S. Ermon, arXiv:2111.08005.

[82] Q. Zhuang and Z. Zhang, Phys. Rev. X **9**, 041023 (2019).

[83] Y. Xia, W. Li, Q. Zhuang, and Z. Zhang, Phys. Rev. X **11**, 021047 (2021).