ONE-STEP PURE-QUANTUM STATE GENERATION VIA CLASSICAL NOISE RE-UPLOADING

Anonymous authors
Paper under double-blind review

ABSTRACT

We propose a novel pure-quantum generative model paradigm that fundamentally avoids the issue of extremely small post-selection probabilities present in previous models. Unlike existing methods that require multi-step noise, addition denoising and classical neural networks, this paradigm enables direct single-step generation of quantum data, significantly improving generation efficiency while substantially reducing the complexity of training and quantum state preparation. Furthermore, by directly sampling classical noise to generate quantum states, the sampling process becomes easier to implement. Experimental results demonstrate that this paradigm outperforms existing quantum generative models in terms of generation quality.

1 INTRODUCTION

The goal of quantum state generation is to learn a target distribution from a small collection of quantum states and then prepare new states from that distribution. QuDDPM addresses this task directly by learning a multi-step map between distributions of individual quantum states (Zhang et al., 2024; Kwun et al., 2025). QFM studies a related but different objective: its trainable tasks supervise application-dependent physical quantities rather than a general distance between complete state distributions (Cui et al., 2026). These developments raise a simpler architectural question: can an easy-to-sample classical random vector be mapped directly to a target state distribution by one quantum circuit evaluation without post-selection?

The main difficulty is that a deterministic quantum channel is linear in the input density operator, whereas the desired map between two state distributions is generally nonlinear. Existing diffusion and flow constructions obtain flexible ensemble maps by combining parameterized circuits with multiple intermediate stages and, when needed, ancilla measurements. We instead avoid learning a quantum-to-quantum transport map: the generator starts from a fixed reference state, while a classical noise vector selects the output through data encoding. This change removes the need to prepare intermediate quantum ensembles during generation.

This observation suggests a different source of nonlinearity: vary a classical input instead of relying on intermediate quantum-state transformations. Figure 1 shows our resulting generator. A noise vector $\boldsymbol{x}$ is repeatedly encoded into a parameterized quantum circuit, and one forward circuit evaluation maps $\boldsymbol{x}$ directly to an output state. Classical sampling supplies the randomness, so the quantum circuit is deterministic once $\boldsymbol{x}$ is fixed. The model consequently needs neither intermediate quantum states nor measured ancillas. Sampling classical noise reduces input preparation costs, while one-step generation lowers both training and generation costs. Eliminating post-selection further reduces generation costs by ensuring a success probability of one. Our method achieves these cost reductions across multiple stages while also delivering better performance.

Our contributions are:

1. **Direct quantum distribution learning.** We introduce a one-step generator based on classical noise re-uploading, with no learned classical decoder, ancillary measurements, or post-selection during generation. Disjoint latent intervals enable conditional sampling with shared circuit parameters.

Reviewers: please read the Reviewer Guidelines (iclr.cc/Conferences/2027/ReviewerGuidelines) and the AI Policy for Reviewers (iclr.cc/Conferences/2027/AIPolicyForReviewers). If you used AI to expand, edit, or polish your review, please provide the input text to the LLM. Better still, consider skipping the LLM and submitting your original text: we, and the authors, are much more interested in your unedited thoughts than in what an LLM has to say. AI-assisted or not, you are putting your name and reputation behind your review: LLM-generated falsehoods, hallucinations or misrepresentations are subject to disciplinary action, which may include desk-rejecting all papers you have authored.

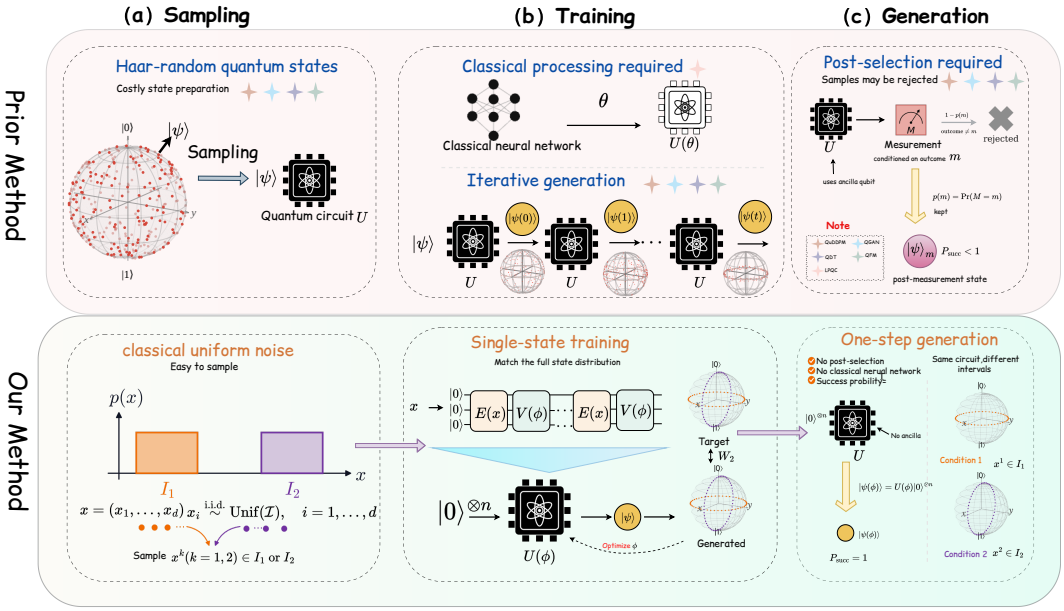


Figure 1: Comparison of prior quantum generative approaches with our one-step generator. (a) Prior methods sample Haar-random quantum states, whereas our method samples classical noise, including from disjoint intervals for conditional generation. (b) Prior approaches may require classical neural processing or iterative quantum evolution; our method directly trains a single circuit to match the target state distribution. (c) Prior methods may rely on post-selection, while our circuit produces each state in one forward evaluation with success probability one and no ancilla measurement.

2. **Distributional expressivity guarantees.** We prove local 2-Wasserstein universality for pure-state distributions under explicit control assumptions, using forward encoders and parameter layers whose individual depth is independent of the re-uploading number.
3. **Empirical accuracy and resource benefits.** Our pure-quantum method improves distribution fitting while reducing generation to a single circuit evaluation with unit ideal success probability. It also achieves more accurate recovery of the evaluated physical statistics with substantially fewer trainable parameters.

2 Related Work

Pure-quantum generative models. Quantum generative adversarial networks formulate quantum-data generation as an adversarial game (Lloyd & Weedbrook, 2018; Ding et al., 2025). Quantum direct transport (QDT), used as a baseline in the QuDDPM study, maps Haar-random input states to a target state distribution with a trainable quantum circuit (Zhang et al., 2024). QuDDPM learns distributions of individual quantum states through a forward scrambling process and measurement-assisted reverse steps (Zhang et al., 2024); its mixed-state extension uses depolarizing noise (Kwun et al., 2025). Chaotic quantum diffusion generates projected ensembles through chaotic Hamiltonian evolution (Tran et al., 2026a). QFM's demonstrated learning tasks target physical criteria such as energy, magnetization, and entanglement entropy (Cui et al., 2026). Many-body projected ensembles (MPEs) use ancilla measurements to generate pure-state ensembles and admit distributional universality guarantees (Tran et al., 2026c). Our method directly fits state distributions in one circuit evaluation, without forward diffusion, intermediate ensembles, measured ancillas, or post-selection.

Classical–quantum generative models. Variational quantum generators encode classical noise and return observable expectation values as classical samples (Romero & Aspuru-Guzik, 2021). (Barthe et al., 2025) prove universality for these classical output distributions, including with re-uploading architectures. Our target is a distribution of quantum states. LPQC uses classical neural

networks to produce data-dependent circuit angles and invokes classical universal approximation in its expressivity proof (Tran et al., 2026b). Our model has no learnable classical decoder: $\boldsymbol{x}$ enters through repeated encodings, with quantum controls fixed after training. Here, "pure-quantum" describes this architecture, not output-state purity.

Data re-uploading and quantum function approximation. Data re-uploading interleaves repeated input encodings with trainable gates. It was introduced for small quantum classifiers and later studied as a function approximator (Pérez-Salinas et al., 2020; 2021). Encoding gates determine the accessible Fourier frequencies, and repeated encodings can expand the spectrum (Schuld et al., 2021). Depth alone, however, need not improve high-dimensional prediction in narrow models (Wang et al., 2025). In contrast to these pointwise or supervised settings, we use re-uploading to push low-dimensional noise to a probability distribution over quantum states.

Table 1 summarizes the architectural distinctions most relevant to our setting. Here, "one-step" means that generation does not require an iterative diffusion or optimization trajectory, and "fits quantum distributions" means learning a distribution of individual quantum states rather than only generating states that satisfy a prescribed physical criterion. "Samples classical noise" means that the latent randomness is sampled as classical variables rather than prepared as random quantum states.

Table 1: Comparison of representative quantum generative methods. A check mark indicates that the method has the corresponding property. In the "No post-selection" column, a check mark means that the ideal generation success probability is one.

Method	Samples classical noise	No post-selection	One-step generation	Fits quantum distributions	No classical neural network
QuDDPM	✗	✗	✗	✓	✓
QDT	✗	✗	✓	✓	✓
QGAN	✗	✗	✓	✓	✓
QFM	✗	✗	✗	✗	✓
LPQC	✓	✓	✓	✓	✗
Ours	✓	✓	✓	✓	✓

3 METHOD

Our generator maps classical noise directly to quantum states, while the training objective matches their distribution to the target. We present the pure-state implementation and its expressivity guarantee below; Appendix F discusses the extension to mixed states.

3.1 ONE-STEP GENERATOR

Let $\boldsymbol{x} = (x_1, x_2, x_3) \sim \mu$ be a three-dimensional noise vector with an easy-to-sample distribution μ. For a one-qubit angle vector $\boldsymbol{\phi} = (\phi_1, \phi_2, \phi_3)$, define

$$\mathrm{Rot}(\boldsymbol{\phi}) = R_z(\phi_3) R_y(\phi_2) R_z(\phi_1).$$

For n qubits, the fixed noise encoder is

$$E_n(\boldsymbol{x}) = e(\boldsymbol{x})^{\otimes n}, \qquad e(\boldsymbol{x}) = \mathrm{Rot}(\boldsymbol{x}). \tag{1}$$

Between encodings, a trainable, data-independent parameter layer $V_p(\boldsymbol{\theta}_p)$ applies at most one rotation–entangling sublayer. The complete generator is

$$U_{P,\Theta}(\boldsymbol{x}) = V_P(\boldsymbol{\theta}_P) E_n(\boldsymbol{x}) V_{P-1}(\boldsymbol{\theta}_{P-1}) \cdots V_1(\boldsymbol{\theta}_1) E_n(\boldsymbol{x}), \tag{2}$$

where P is the number of times the same $\boldsymbol{x}$ is encoded and $\Theta = (\boldsymbol{\theta}_1, \ldots, \boldsymbol{\theta}_P)$ contains all trainable control parameters. Increasing P enriches the dependence on $\boldsymbol{x}$. Given $\boldsymbol{x}$, the generated pure state and its density operator are

$$|\psi_\Theta(\boldsymbol{x})\rangle = U_{P,\Theta}(\boldsymbol{x})|0^n\rangle, \qquad \rho_\Theta(\boldsymbol{x}) = |\psi_\Theta(\boldsymbol{x})\rangle\langle\psi_\Theta(\boldsymbol{x})|.$$

One draw of $\boldsymbol{x}$ followed by one evaluation of Eq. (2) produces one sample. No measurement outcome is accepted or rejected, so the ideal success probability is one.

Figure 2 illustrates the complete generator and a representative realization of one parameter layer.

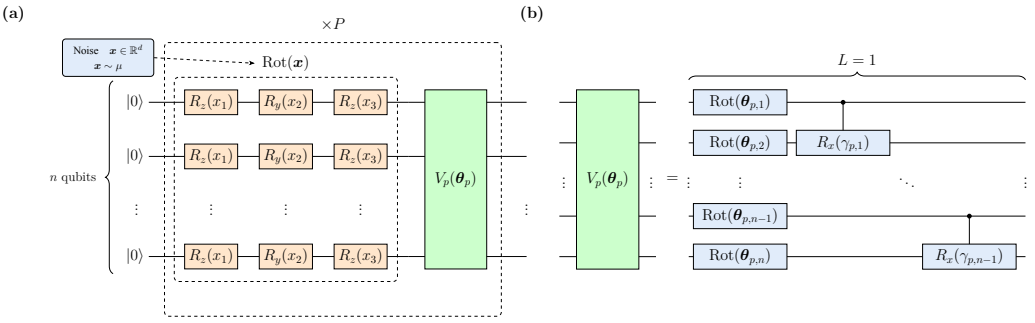


Figure 2: Circuit architecture of the one-step generator. (a) The same classical noise vector $\boldsymbol{x}$ is encoded as $\mathrm{Rot}(\boldsymbol{x})$ on every qubit, followed by a trainable parameter layer $V_p(\boldsymbol{\theta}_p)$; this pair is repeated P times. (b) A representative parameter layer applies independent trainable Rot gates followed by adjacent C-R_x gates, each with one trainable angle. Here $L = 1$ denotes one rotation–entangling sublayer inside each parameter layer, whereas P is the number of noise re-uploadings.

Some experiments use trainable positive scales, $e_{\boldsymbol{s}}(\boldsymbol{x}) = \mathrm{Rot}(s_1x_1, s_2x_2, s_3x_3)$, while one-dimensional instances tie the coordinates as $\boldsymbol{x} = (x, x, x)$. These are practical specializations of Eq. (1).

3.2 Distribution Matching and Conditional Generation

We train the generator by comparing finite batches from the target and generated distributions. Let $\{|\phi_i\rangle\}_{i=1}^M$ be target states and let $\boldsymbol{x}^{(1)}, \ldots, \boldsymbol{x}^{(N)} \sim \mu$ be noise samples. The generated batch is $\{|\psi_\Theta(\boldsymbol{x}^{(j)})\rangle\}_{j=1}^N$. Its pairwise cost matrix with the target batch is

$$C_{ij} = 1 - |\langle\phi_i|\psi_\Theta(\boldsymbol{x}^{(j)})\rangle|^2, \tag{3}$$

To make this cost explicit, consider two pure states $\rho = |\psi\rangle\langle\psi|$ and $\sigma = |\phi\rangle\langle\phi|$. Their trace distance is

$$d(\rho, \sigma) = \tfrac{1}{2}\|\rho - \sigma\|_1 = \sqrt{1 - |\langle\psi|\phi\rangle|^2}. \tag{4}$$

Therefore, the pairwise cost in Eq. (3) is exactly the squared trace distance. The overlaps between the target and generated states can be estimated with a SWAP test (Barenco et al., 1997). Following (Zhang et al., 2024), we minimize the empirical squared 2-Wasserstein cost

$$W_2^2 = \min_{\Pi \geq 0} \sum_{i=1}^{M} \sum_{j=1}^{N} \Pi_{ij} C_{ij}, \quad \text{s.t.} \quad \Pi \mathbf{1}_N = \mathbf{a}, \quad \Pi^\top \mathbf{1}_M = \mathbf{b}. \tag{5}$$

Here Π is a transport plan and $\mathbf{a}, \mathbf{b}$ are the empirical sample weights (uniform in our experiments) (Peyré & Cuturi, 2019). We solve the entropy-regularized problem with the Sinkhorn algorithm (Cuturi, 2013). Minimizing W_2^2 and W_2 has the same minimizers.

Conditional generation uses latent support itself as the label, rather than appending a one-hot input. We assign disjoint intervals $I_1, \ldots, I_K$ to K target categories and train one circuit across all categories. For the three-dimensional encoder used in the conditional experiments, write $\boldsymbol{z} = (z_1, z_2, z_3)$ and sample

$$\mathbb{P}(\boldsymbol{z} \mid c) = \prod_{k=1}^{3} \mathbb{P}(z_k \mid c), \quad \text{where} \quad \mathbb{P}(z_k \mid c) = \mathrm{Unif}(I_c), \quad k = 1, 2, 3, \tag{6}$$

so category c occupies the cube I_c^3. At inference time, sampling from I_c^3 requests category c while all circuit parameters remain fixed. Conditioning therefore requires neither an ancillary qubit nor measurement-based post-selection.

3.3 EXPRESSIVITY: ASSUMPTIONS AND MAIN RESULTS

Function or amplitude approximation alone does not show that a generator can match an entire distribution of quantum states (Pérez-Salinas et al., 2020; 2021). We therefore prove a distribution-level result for the architecture in Eq. (2). The proofs are in Appendix A.

Let $D = 2^n$ and define $\mathcal{S}_n = \{\rho = |\psi\rangle\langle\psi| : |\psi\rangle \in (\mathbb{C}^2)^{\otimes n},\ \langle\psi|\psi\rangle = 1\}$, the set of n-qubit pure-state density operators. In Theorem 1, W_2 is the distance between probability distributions on $\mathcal{S}_n$ whose ground metric is the trace distance in Eq. (4). Sampling $\boldsymbol{x} \sim \mu$ produces the random state $\rho_\Theta(\boldsymbol{x})$; its distribution is denoted by $Q_{P,\Theta}$.

Let $\mathcal{V} \subset \mathrm{SU}(D)$ be the set of data-independent unitaries that one fixed parameter layer can realize. A parameter layer may be omitted, so $I_D \in \mathcal{V}$. Write

$$\langle\mathcal{V}\rangle_+ = \{A_m \cdots A_1 : m \geq 0,\ A_j \in \mathcal{V}\}$$

for the set of all finite circuits obtained by concatenating these parameter layers. More explicitly, the case $m = 0$ is the empty product I_D, the case $m = 1$ gives one parameter layer from $\mathcal{V}$, and $m = 2$ gives a product A_2A_1 of two allowed layers; arbitrary finite values of m are allowed. The subscript "+" means that these are forward products only: the notation does not insert a formal inverse A_j^{-1} unless that inverse itself is already available as another allowed parameter setting in $\mathcal{V}$. Thus $\langle\mathcal{V}\rangle_+$ is a semigroup (closed under concatenation), not necessarily a group. In the theorem, every $V_p(\boldsymbol{\theta}_p)$ in Eq. (2) is selected from the same set $\mathcal{V}$, including its allowed bypass setting.

Assumption 1 (Realizable parameter layers of fixed depth)**.** *Fix n and $D = 2^n$. The architecture and depth of one parameter layer are fixed independently of the re-uploading number P. The set $\mathcal{V}$ contains all data-independent unitaries obtainable by changing the parameters of this one parameter layer. We assume the following:*

1. *The parameter layer can be omitted (bypassed), so $I_D \in \mathcal{V}$.*
2. *By concatenating finitely many allowed parameter layers, we can approximate any n-qubit special unitary. Equivalently, the positive control semigroup is dense:*
$$\overline{\langle\mathcal{V}\rangle_+} = \mathrm{SU}(D), \tag{7}$$
3. *Every phase-adjusted n-qubit Pauli operation used in the twirling construction can be realized within one parameter layer.*

Condition 3 means that, for every required tensor-product Pauli operation

$$W = W_1 \otimes \cdots \otimes W_n, \qquad W_q \in \{\pm I_2, \pm \mathrm{i}X, \pm \mathrm{i}Y, \pm \mathrm{i}Z\},$$

there is a parameter setting $\boldsymbol{\theta}_W$ such that $V(\boldsymbol{\theta}_W) = W$. At these settings, the full parameter layer reduces to a tensor product of single-qubit Rot gates. In our concrete construction, each entangling gate is C-$R_x(\gamma)$ with one trainable angle. It can be turned off by setting $\gamma = 0$, for which C-$R_x(0) = I_4$. At other parameter settings, the same layer can still generate entanglement. Placing these special parameter layers between consecutive forward encoders cancels the leading error; repeating the construction ultimately allows the distribution-fitting error to be made arbitrarily small. The detailed argument is given in Appendix A.

Assumption 1 is a sufficient condition, not a claimed necessary condition. It does not require one parameter layer to implement an arbitrary unitary: universality is required only from finite products of the same fixed-depth layers.

Following the Fourier perspective on data encoding (Schuld et al., 2021), we bound the matrix-element frequencies for our specific architecture.

Lemma 1 (Fourier order grows with re-uploading depth)**.** *Let $z_k = e^{\mathrm{i}x_k/2}$ for $k = 1, 2, 3$. Every matrix entry of $U_{P,\Theta}(\boldsymbol{x})$ is a multivariate Laurent polynomial in (z_1, z_2, z_3), and every exponent m_k satisfies $|m_k| \leq nP$. A valid choice of controls attains the simultaneous extreme order $|m_1| = |m_2| = |m_3| = nP$.*

Lemma 1 isolates the role of P: increasing the number of re-uploadings permits progressively higher Fourier orders in all three coordinates. The next result establishes distributional universality without inserting a reverse encoding block.

Theorem 1 (Local forward-only universality for pure-state distributions)**.** *Fix a number of qubits $n \geq 1$ and a P-independent set of realizable parameter layers $\mathcal{V}$ satisfying Assumption 1. Then there exists $a_0(n) > 0$ with the following property. Let*

$$K_a = [-a, a]^3, \qquad 0 < a < a_0(n),$$

and let μ be a probability measure supported on K_a with a positive continuous density there. For every probability distribution Q on $\mathcal{S}_n$ and every $\varepsilon > 0$, there are a finite P and parameters Θ such that

$$W_2\big(Q, Q_{P,\Theta}\big) < \varepsilon.$$

Every occurrence of the data encoder in $U_{P,\Theta}$ is the forward encoder $E_n(\boldsymbol{x})$ from Eq. (1)*; $E_n(\boldsymbol{x})^{-1}$ is never added to the circuit architecture.*

Proof idea. The special parameter settings in Condition 3 cancel the leading error when they are inserted between consecutive forward encoders. Repeating this construction lets a forward-only circuit reproduce the effect of the inverse encoder. Together with the control ability in Condition 2, the resulting circuits can approximate the continuous state maps needed to assign arbitrary probability masses to target states. The distribution-fitting error can therefore be made arbitrarily small. The complete proof is given in Appendix A.

Relation to prior universality results. Universality has been established for classical expectation-value distributions (Barthe et al., 2025), measurement-generated pure-state ensembles (Tran et al., 2026c), and latent-conditioned circuits with classical neural decoders (Tran et al., 2026b). For the trace-distance ground metric, $d \leq 1$ gives $W_1 \leq W_2 \leq \sqrt{W_1}$, so W_1 universality also implies W_2 universality. Our result concerns the architecture: three-dimensional noise on a sufficiently small cube, only forward encoders, and neither measured ancillas nor a classical decoder, under Assumption 1. Its expressivity follows from Pauli twirling, Lie closure, and Stone–Weierstrass.

Theorem 1 shows that suitable circuit parameters exist, but does not guarantee that training will find them or that only a few re-uploadings are needed. It also does not imply that adding more trainable layers per re-uploading improves training.

4 Experiments

We evaluate our method on multiple quantum-state datasets, including the 0-cluster, the $\langle Y \rangle^2 = 0$ circle, correlated-noise states, and ten-qubit transverse-field Ising ground states. We also investigate entanglement-entropy fitting and conditional generation. Comparisons include QuDDPM, QFM, QDT, QGAN, and LPQC on the corresponding tasks. Detailed experimental settings are provided in Appendix D.1, and ablation studies are reported in Appendix E.

4.1 One-Step Generation without Post-Selection

We compare our method with QuDDPM on its single-qubit 0-cluster dataset (Zhang et al., 2024), a finite-width distribution around $|0\rangle$:

$$|\phi_k\rangle = \frac{|0\rangle + 0.08(a_k + ib_k)|1\rangle}{\sqrt{1 + 0.08^2(a_k^2 + b_k^2)}}, \qquad a_k, b_k \overset{\text{i.i.d.}}{\sim} \mathcal{N}(0, 1).$$

We use 100 training states and 1000 independent held-out target states.

QuDDPM transforms Haar-random states through T reverse denoising steps, each using four parameter layers and ancilla post-selection. We vary $T = 1, \ldots, 20$ and set $P = 4T$, matching the $4T$ trainable Rot gates and C-R_x entangling gates on the data qubit; QuDDPM additionally trains ancilla gates. Both methods use the loss in Eq. (5). Architecture and training settings are given in Appendix D.1.1.

Figure 3 compares the empirical Wasserstein cost between generated quantum-state ensembles and held-out target sets for different re-uploading numbers P in our model and denoising steps T in QuDDPM, with $P = 4T$. Our model achieves a lower cost across all tested settings. QuDDPM's end-to-end success probability decreases with T, sometimes falling below 10^{-9}, while ours remains one. For batches of M target and N generated states, our loss requires MN pairwise overlap estimates, compared with TMN across QuDDPM's denoising objectives.

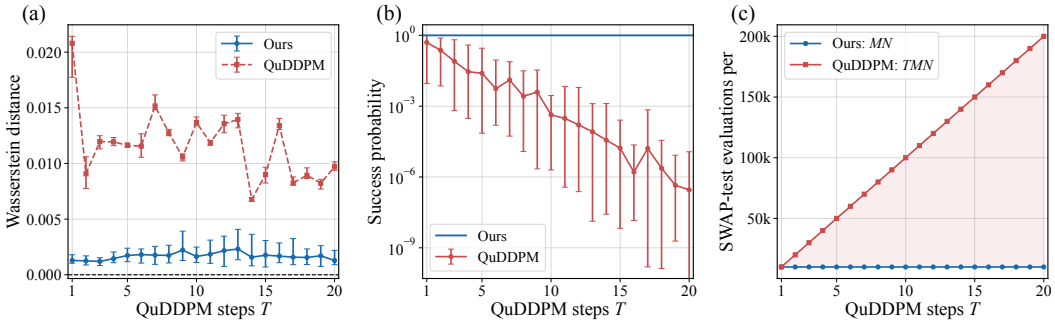


Figure 3: Comparison with QuDDPM on the 0-cluster dataset for $T = 1, \ldots, 20$. We set $P = 4T$ to match trainable gate. In panels (a) and (b), markers show means over five independent runs, and error bars span the observed minimum and maximum. (a) Empirical Wasserstein objective between 1000 generated states and 1000 held-out target states; lower is better. (b) End-to-end success probability of the selected QuDDPM post-selection branch; our method has success probability one. (c) Pairwise SWAP-test evaluations per complete loss for $M = N = 100$: MN for our method and TMN for QuDDPM.

4.2 Improved Quantum State Distribution Fitting Performance

We compare with QDT (Zhang et al., 2024) and QGAN (Lloyd & Weedbrook, 2018; Ding et al., 2025) on the real-amplitude circle

$$|\phi(\varphi)\rangle = \cos(\varphi)|0\rangle + \sin(\varphi)|1\rangle, \qquad \varphi \sim \mathrm{Unif}[0, 2\pi). \tag{8}$$

whose Bloch vectors form the x–z great circle and satisfy $\langle Y \rangle^2 = 0$. All methods use 100 training states, 2000 training epochs, and five independent initializations; our generator has $P = 20$. QDT uses a fidelity-kernel distribution loss and QGAN an adversarial loss, defined in Appendix B. Full configurations are in Appendix D.1.2.

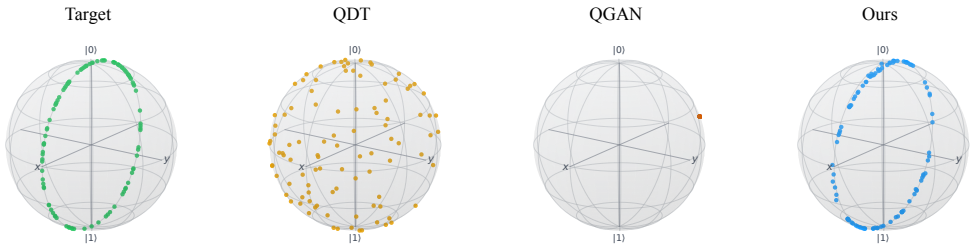


Figure 4: Generated-state comparison on the $\langle Y \rangle^2 = 0$ circle. From left to right: the target distribution, QDT, QGAN, and our generator. Each panel displays 100 states. Under the common oblique Bloch-sphere view, the ideal x–z great circle appears as a single vertical ellipse. QDT leaves the target plane, QGAN collapses to one mode, and our samples recover nearly the complete circle.

Figure 4 displays representative generated samples; its saved run configurations are detailed in Appendix D.1.2. QDT produces states outside the target plane, QGAN collapses near one state, and our generator recovers nearly the full target circle.

4.3 Improved Entanglement Fitting

Following QFM (Cui et al., 2026), we train a two-qubit generator to match prescribed entanglement entropies $S(\rho_1) = -\mathrm{Tr}(\rho_1 \log_2 \rho_1)$, where ρ_1 is the first qubit's reduced state. We fit 11 targets $S_{\mathrm{tar}} = 0, 0.1, \ldots, 1$ with 20 QFM layers or $P = 20$ re-uploadings and report five independent runs. The entropy loss and training settings are in Appendix D.1.3.

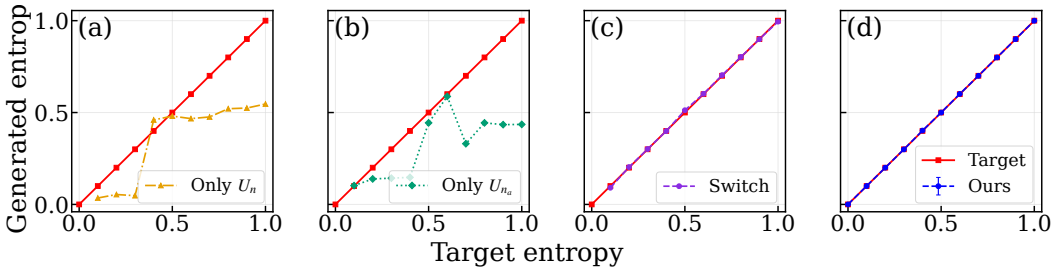


Figure 5: Entanglement-entropy fitting for QFM and our generator. (a) QFM using U_n on the data qubits only; (b) QFM using the ancilla-assisted U_{n_a}; (c) QFM dynamically switching between U_n and U_{n_a}; and (d) our method. Red squares show target entropy. Blue circles show the mean generated entropy over five independent runs, and error bars show the standard deviation across runs.

Figure 5 shows visible deviations for the two fixed QFM modes at several target values. Dynamic switching improves QFM's fit but requires changing between a data-only circuit and an ancilla-assisted circuit across target steps. Our method follows all 11 target values at the plotted resolution with one fixed circuit form and no ancillary-qubit switching. While QFM is designed for criterion-conditioned state generation, our method also accurately fits the prescribed entanglement entropies, in addition to learning distributions of quantum states.

4.4 Quantum Phase Noise Estimation

We adopt the two-qubit correlated coherent gate-error dataset from the QuDDPM study (Zhang et al., 2024) and compare our generator with LPQC (Tran et al., 2026b) on estimating the XX-error probability $p = 0.8$. The mean $|10\rangle$ population of generated states provides an estimate $\widehat{p}$; Appendix C gives the noise model and measurement derivation. Both methods use the same 200 training states, and we evaluate $N \in \{1000, 2000, 5000\}$ generated states per run. Training details are in Appendix D.1.4.

For each generated-sample size N, let $\widehat{p}_r$ denote the mixing-weight estimate obtained in run r. We first compute the absolute error $|\widehat{p}_r - p|$ relative to the true value $p = 0.8$, then report $\mathrm{mean}_r\,|\widehat{p}_r - p| \pm \mathrm{std}_r(|\widehat{p}_r - p|)$ over five independent runs. These statistics summarize the typical estimation error and its variability across runs. The absolute value is taken before averaging, preventing overestimation and underestimation from cancelling each other. This differs from $|\,\mathrm{mean}_r(\widehat{p}_r) - p|$, which measures the error of the averaged estimate and can conceal larger errors in individual runs.

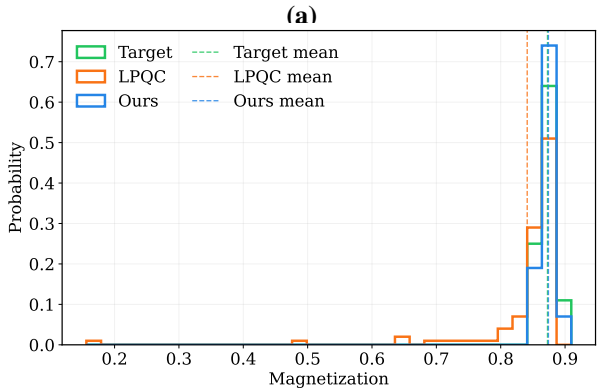
Table 2: Absolute error in estimating the correlated-noise mixing weight $p = 0.8$ at different generated-sample sizes. Values are the mean $\pm$ standard deviation over five independent runs; lower is better.

Method	Parameters	Error $\lvert\widehat{p} - p\rvert$ ($N = 1000$)	Error $\lvert\widehat{p} - p\rvert$ ($N = 2000$)	Error $\lvert\widehat{p} - p\rvert$ ($N = 5000$)
LPQC	5164	0.0488 ± 0.0196	0.0476 ± 0.0232	0.0391 ± 0.0343
Ours	**143**	$\mathbf{0.0205 \pm 0.0218}$	$\mathbf{0.0215 \pm 0.0182}$	$\mathbf{0.0215 \pm 0.0174}$

Table 2 shows that our generator achieves lower mean absolute error than LPQC across all tested sample sizes. This improvement is obtained with substantially fewer trainable parameters, indicating that accurate recovery of the correlated-noise mixing weight can be achieved with a more compact model.

4.5 Quantum Phase Learning

We learn a distribution of ground states of the ten-qubit transverse-field Ising model (TFIM), using 100 training states and 100 new generated states for evaluation. We compare their transverse magnetization $M_x = n^{-1} \sum_{j=1}^{n} \langle X_j \rangle$; the Hamiltonian and training settings are given in Appendix D.1.5.



(b)

Method	Mean M_x	Std. M_x	Trainable parameters
Target	0.87267	0.01116	–
LPQC	0.84117	0.08822	135,588
Ours	**0.87415**	**0.00940**	**5,703**

Figure 6: Quantum phase learning on the ten-qubit TFIM. (a) Empirical probability distributions of transverse magnetization for the target states, LPQC, and our generator; dashed lines indicate the corresponding means. (b) Magnetization statistics and trainable parameter counts. Bold highlights the better generative-model result in each column.

Figure 6 shows that our magnetization distribution closely follows the target, while LPQC's is substantially broader. With substantially fewer trainable parameters than LPQC, our model better reproduces both the target mean and spread. This agreement demonstrates accurate recovery of the tested magnetization statistics.

4.6 Conditional Generation

We use disjoint latent intervals to select target distributions through Eq. (6). One model learns the $\langle Z\rangle^2 = 0$ and $\langle Y\rangle^2 = 0$ circles; a separately trained model learns four finite-width clusters around $|0\rangle$, $|1\rangle$, $|+\rangle$, and $|-\rangle$. Each model uses one qubit, $P = 20$, with 200 training and 1000 test states per condition. Within each model, conditions share all circuit parameters and latent scales; generation changes only the sampling interval. Appendix D.1.6 specifies the intervals and training settings.

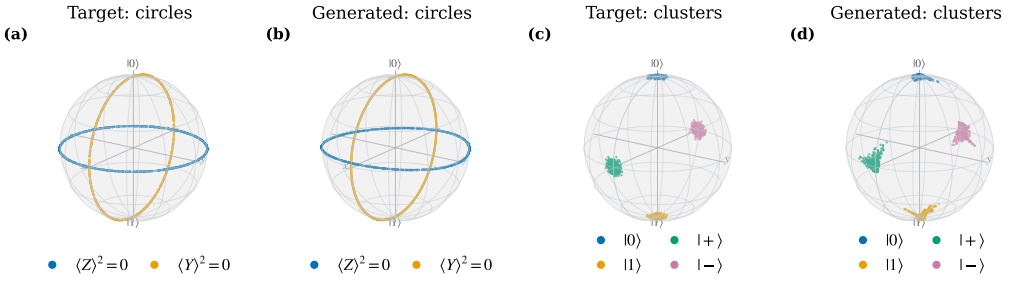


Figure 7: Conditional generation. Target and generated samples for (a,b) two great circles and (c,d) four finite-width clusters. Each condition contains 1000 states; the circle and cluster models are trained separately.

Figure 7 shows that the generated modes are clearly distinguishable, with each mode closely matching its corresponding target distribution. Changing the latent sampling interval reliably selects the intended circle or cluster while preserving its characteristic geometry.

5 Conclusion

We propose classical noise re-uploading, a new paradigm for quantum state generation. It features easy-to-sample noise, one-step generation, no post-selection, and unit success probability, making it a strong candidate for quantum state generation.

AI USE STATEMENT

We used generative AI tools only for language polishing, figure preparation, and assistance with checking mathematical proofs. The authors reviewed and verified all AI-assisted material and take full responsibility for the final content of this work.

REPRODUCIBILITY STATEMENT

Section 3 defines the generator, training objective, and theoretical assumptions. Appendix D.1 records datasets and training configurations. We will release the code after the manuscript is accepted.

REFERENCES

Adriano Barenco, André Berthiaume, David Deutsch, Artur Ekert, Richard Jozsa, and Chiara Macchiavello. Stabilization of quantum computations by symmetrization. *SIAM Journal on Computing*, 26(5):1541–1557, 1997.

Alice Barthe, Michele Grossi, Sofia Vallecorsa, Jordi Tura, and Vedran Dunjko. Parameterized quantum circuits as universal generative models for continuous multivariate distributions. *npj Quantum Information*, 11(1):121, 2025.

Ranyiliu Chen, Zhixin Song, Xuanqiang Zhao, and Xin Wang. Variational quantum algorithms for trace distance and fidelity estimation. *Quantum Science and Technology*, 7(1):015019, 2022.

Zidong Cui, Pan Zhang, and Ying Tang. Quantum flow matching. *PRX Intelligence*, 1(1):013009, 2026.

Marco Cuturi. Sinkhorn distances: Lightspeed computation of optimal transport. In *Advances in Neural Information Processing Systems*, volume 26, pp. 2292–2300. Curran Associates, Inc., 2013.

Yifan Ding, Zi Li, and Nanrun Zhou. Quantum generative adversarial network based on the quantum Born machine. *Advanced Engineering Informatics*, 68:103622, 2025.

Artur K. Ekert, Carolina Moura Alves, Daniel K. L. Oi, Michał Horodecki, Paweł Horodecki, and L. C. Kwek. Direct estimations of linear and nonlinear functionals of a quantum state. *Physical Review Letters*, 88(21):217901, 2002.

Arthur Gretton, Karsten M. Borgwardt, Malte J. Rasch, Bernhard Schölkopf, and Alexander Smola. A kernel two-sample test. *Journal of Machine Learning Research*, 13(25):723–773, 2012.

Jeongwan Haah, Aram W. Harrow, Zhengfeng Ji, Xiaodi Wu, and Nengkun Yu. Sample-optimal tomography of quantum states. *IEEE Transactions on Information Theory*, 63(9):5628–5641, 2017.

Diederik P. Kingma and Jimmy Ba. Adam: A method for stochastic optimization. In *International Conference on Learning Representations*, 2015.

Gino Kwun, Bingzhi Zhang, and Quntao Zhuang. Mixed-state quantum denoising diffusion probabilistic model. *Physical Review A*, 111(3):032610, 2025.

Seth Lloyd and Christian Weedbrook. Quantum generative adversarial learning. *Physical Review Letters*, 121(4):040502, 2018.

Adrián Pérez-Salinas, Alba Cervera-Lierta, Elies Gil-Fuster, and José I. Latorre. Data re-uploading for a universal quantum classifier. *Quantum*, 4:226, 2020.

Adrián Pérez-Salinas, David López-Núñez, Artur García-Sáez, P. Forn-Díaz, and José I. Latorre. One qubit as a universal approximant. *Physical Review A*, 104(1):012405, 2021.

Gabriel Peyré and Marco Cuturi. Computational optimal transport: With applications to data science. *Foundations and Trends in Machine Learning*, 11(5–6):355–607, 2019.

Jonathan Romero and Alán Aspuru-Guzik. Variational quantum generators: Generative adversarial quantum machine learning for continuous distributions. *Advanced Quantum Technologies*, 4(1): 2000003, 2021.

Maria Schuld, Ryan Sweke, and Johannes Jakob Meyer. Effect of data encoding on the expressive power of variational quantum-machine-learning models. *Physical Review A*, 103(3):032430, 2021.

Quoc Hoan Tran, Koki Chinzei, Yasuhiro Endo, and Hirotaka Oshima. Learning quantum data distribution via chaotic quantum diffusion model. arXiv preprint arXiv:2602.22061, 2026a.

Quoc Hoan Tran, Koki Chinzei, Yasuhiro Endo, and Hirotaka Oshima. Latent-conditioned parameterized quantum circuits as universal approximators for distributions over quantum states. arXiv preprint arXiv:2605.28690, 2026b.

Quoc Hoan Tran, Koki Chinzei, Yasuhiro Endo, and Hirotaka Oshima. Universality of many-body projected ensemble for learning quantum data distribution. *arXiv preprint arXiv:2601.18637*, 2026c.

Qisheng Wang and Zhicheng Zhang. Fast quantum algorithms for trace distance estimation. *IEEE Transactions on Information Theory*, 70(4):2720–2733, 2024.

Xin Wang, Hanxiao Tao, and Rebing Wu. Predictive performance of deep quantum data re-uploading models. In *Proceedings of the 42nd International Conference on Machine Learning*, volume 267 of *Proceedings of Machine Learning Research*, pp. 64491–64524. PMLR, 2025.

Bingzhi Zhang, Peng Xu, Xiaohui Chen, and Quntao Zhuang. Generative quantum machine learning via denoising diffusion probabilistic models. *Physical Review Letters*, 132(10):100602, 2024.

A PROOFS FOR THE EXPRESSIVITY RESULTS

We restate each result from the main text before giving its proof. All notation and assumptions are as defined in Section 3.

A.1 FOURIER ORDER UNDER RE-UPLOADING

Lemma A.1 (Lemma 1 in the main text)**.** *Let $z_k = e^{\mathrm{i}x_k/2}$ for $k = 1, 2, 3$. Every matrix entry of $U_{P,\Theta}(\boldsymbol{x})$ is a multivariate Laurent polynomial in (z_1, z_2, z_3), and every exponent m_k satisfies $|m_k| \leq nP$. A valid choice of controls attains the simultaneous extreme order $|m_1| = |m_2| = |m_3| = nP$.*

Proof. Let $z_k = e^{\mathrm{i}x_k/2}$. The one-qubit encoder in Eq. (1) is

$$e(\boldsymbol{x}) = \begin{pmatrix} \frac{1}{2} z_3^{-1} z_1^{-1}(z_2 + z_2^{-1}) & -\frac{1}{2\mathrm{i}} z_3^{-1} z_1 (z_2 - z_2^{-1}) \\ \frac{1}{2\mathrm{i}} z_3 z_1^{-1}(z_2 - z_2^{-1}) & \frac{1}{2} z_3 z_1 (z_2 + z_2^{-1}) \end{pmatrix}.$$

Every entry has exponent $+1$ or -1 in each variable. Therefore every entry of $E_n(\boldsymbol{x}) = e(\boldsymbol{x})^{\otimes n}$ has coordinate-wise Laurent degree at most n. Multiplication by a data-independent control does not change the degree, and each additional encoding increases every coordinate-wise bound by at most n. Induction over the P encodings gives $|m_k| \leq nP$.

The bound is attained by setting all controls to the identity. In the $(0, 0)$ entry of $e(\boldsymbol{x})^P$, the unique multiplication path producing the monomial $z_1^{-P} z_2^{-P} z_3^{-P}$ selects the $z_1^{-1} z_2^{-1} z_3^{-1}/2$ term from the upper-left entry at every factor. Hence the $(0^n, 0^n)$ entry of $E_n(\boldsymbol{x})^P = (e(\boldsymbol{x})^P)^{\otimes n}$ contains

$$2^{-nP} z_1^{-nP} z_2^{-nP} z_3^{-nP},$$

which proves the claimed simultaneous extreme order. □

A.2 LOCAL FORWARD-ONLY DISTRIBUTION UNIVERSALITY

Theorem A.1 (Theorem 1 in the main text)**.** *Fix a number of qubits $n \geq 1$ and a P-independent set of realizable parameter layers $\mathcal{V}$ satisfying Assumption 1. Then there exists $a_0(n) > 0$ with the following property. Let*

$$K_a = [-a, a]^3, \qquad 0 < a < a_0(n),$$

and let μ be a probability measure supported on K_a with a positive continuous density there. For every probability distribution Q on $\mathcal{S}_n$ and every $\varepsilon > 0$, there are a finite P and parameters Θ such that

$$W_2\big(Q, Q_{P,\Theta}\big) < \varepsilon.$$

Every occurrence of the data encoder in $U_{P,\Theta}$ is the forward encoder $E_n(\boldsymbol{x})$ from Eq. (1)*; $E_n(\boldsymbol{x})^{-1}$ is never added to the circuit architecture.*

Proof. We prove the theorem in four steps. Throughout, norms without a subscript are operator norms, and all function-space closures use the uniform norm $\|G\|_\infty := \sup_z \|G(z)\|$.

Step 1: forward words locally synthesize the inverse. For a compact space Z and the signal function $F(z) = z \in \mathrm{SU}(D)$, define the parameter-layer positive words

$$\mathcal{W}_+^{\mathcal{V}}(F) = \{A_P F(\cdot) A_{P-1} F(\cdot) \cdots A_1 F(\cdot) : P \geq 1,\ A_j \in \mathcal{V}\}.$$

Every occurrence of F is forward, and every A_j is one allowed parameter layer. Let $\mathcal{G}_+^{\mathcal{V}}(F)$ be the uniform closure of $\mathcal{W}_+^{\mathcal{V}}(F)$. Pointwise multiplication concatenates two positive words, so $\mathcal{G}_+^{\mathcal{V}}(F)$ is a closed semigroup.

Let

$$\mathcal{Q} = \{\pm I_2, \pm\mathrm{i}X, \pm\mathrm{i}Y, \pm\mathrm{i}Z\}, \qquad \mathcal{H}_n = \{W_1 \otimes \cdots \otimes W_n : W_q \in \mathcal{Q}\}.$$

For completeness, consider a standard multi-qubit parameter layer with independent single-qubit $\mathrm{Rot} = R_z R_y R_z$ gates and parameterized C-R_x gates. Each controlled gate has one scalar angle γ and acts as

$$\text{C-}R_x(\gamma) = |0\rangle\langle 0| \otimes I_2 + |1\rangle\langle 1| \otimes R_x(\gamma).$$

It can be set to the identity through C-$R_x(0) = I_4$. The one-qubit choices in $\mathcal{Q}$ are available because, up to the allowed signs, $R_y(\pi) = -\mathrm{i}Y$, $R_z(\pi) = -\mathrm{i}Z$, and $R_z(\pi)R_y(\pi) = \mathrm{i}X$, while the identity is obtained by setting all angles to zero. For any $W = W_1 \otimes \cdots \otimes W_n \in \mathcal{H}_n$, choose the single-qubit Rot angles to realize each W_q and set every C-R_x gate to the identity. This gives $W \in \mathcal{V}$ in one parameter layer and verifies the third item of Assumption 1. If the entangling gates are mandatory or the local parameters are tied across qubits, this verification does not apply and $\mathcal{H}_n \subseteq \mathcal{V}$ must be checked separately.

The set $\mathcal{H}_n$ defined above is a finite subgroup, and its conjugation twirl satisfies

$$\sum_{W \in \mathcal{H}_n} W A W^\dagger = 0, \qquad A \in \mathfrak{su}(D). \tag{9}$$

Set $K = 3|\mathcal{H}_n|$. Form a list $W_1, \ldots, W_K$ containing every element of $\mathcal{H}_n$ exactly three times, ordered so that $W_1 = I_D$ and $W_{K-1} = W_K = I_D$. Define

$$\Phi(U) = \prod_{j=1}^{K} W_j U W_j^\dagger.$$

The factors are multiplied in the displayed order from $j = 1$ to K. Because $W_K = I_D$, expansion gives

$$\Phi(U) = W_1 U (W_1^\dagger W_2) U \cdots (W_{K-1}^\dagger W_K) U. \tag{10}$$

Here W_1 and every transition $W_j^\dagger W_{j+1}$ lie in $\mathcal{H}_n \subseteq \mathcal{V}$. Thus Eq. (10) uses only one allowed parameter layer between consecutive forward occurrences of U.

In a sufficiently small neighborhood of the origin in $\mathfrak{su}(D)$, the map

$$\phi(A) = \log\big(\Phi(e^A)\big)$$

is analytic. It satisfies $\phi(0) = 0$, and

$$\mathrm{d}\phi|_0[A] = \sum_{j=1}^{K} W_j A W_j^\dagger = 3 \sum_{W \in \mathcal{H}_n} W A W^\dagger = 0$$

by Eq. (9). Finite-dimensional Taylor expansion therefore supplies constants $c_D, s_D > 0$ such that

$$\|\phi(A)\| \le c_D \|A\|^2 \quad \text{whenever} \quad \|A\| \le s_D. \tag{11}$$

Fix

$$0 < r_D < \min\{s_D, (2c_D)^{-1}\}.$$

Suppose $Z \subset \mathrm{SU}(D)$ is compact and $\sup_{z \in Z} \|\log z\| < r_D$. Set $U_0(z) = z$ and $U_{m+1}(z) = \Phi(U_m(z))$. If $U_m(z) = e^{A_m(z)}$, then Eq. (11) gives

$$\eta_{m+1} \le c_D \eta_m^2, \qquad \eta_m \le c_D^{-1} (c_D \eta_0)^{2^m}, \qquad \eta_m := \sup_{z \in Z} \|A_m(z)\|.$$

Thus $U_m \to I_D$ uniformly on Z.

Every U_m is a legal parameter-layer positive word in z whose first control is the identity. This follows by induction. The base word is $U_0(z) = I_D z$. If the claim holds for U_m, substituting it into Eq. (10) preserves the form because all transition controls are in $\mathcal{V}$ and the first control of each substituted copy is I_D. Because the final two list entries are identities, the expanded recursion ends in $U_m U_m$. Thus U_1 ends in zz, and the terminal zz property follows for every later U_m by induction. Hence, for every $m \ge 1$, deleting the final occurrence of z leaves a legal parameter-layer positive word B_m satisfying

$$U_m(z) = B_m(z) z, \qquad \|B_m - z^{-1}\|_\infty = \|U_m - I_D\|_\infty \longrightarrow 0.$$

The inverse signal is therefore in $\mathcal{G}_+^{\mathcal{V}}(F)$ without ever using a reverse query. Because the first control of U_m is I_D, replacing it by any $A \in \mathcal{V}$ gives another legal word AU_m, and $AU_m \to A$ uniformly. Thus every constant from $\mathcal{V}$ belongs to the closure. Closedness under pointwise multiplication first places every constant in $\langle \mathcal{V} \rangle_+$ in the closure. Equation (7) and closedness of $\mathcal{G}_+^{\mathcal{V}}(F)$ then place every constant map in $\mathrm{SU}(D)$ in the closure. Explicitly, if $w(z) = A_P z \cdots A_1 z$, then

$$w(z)^{-1} = z^{-1} A_1^{-1} \cdots z^{-1} A_P^{-1} \in \mathcal{G}_+^{\mathcal{V}}(F),$$

because z^{-1} and all constant unitaries are already in the closed semigroup. Continuity of inversion extends this conclusion to the closure; hence $\mathcal{G}_+^{\mathcal{V}}(F)$ is a group.

Step 2: the closure contains the required continuous unitary maps. Assume in addition that Z has no distinct projectively equivalent points:

$$z, w \in Z, \quad z = \zeta w, \quad \zeta^D = 1 \quad \Longrightarrow \quad z = w. \tag{12}$$

Let $\mathfrak{g} = \mathfrak{su}(D)$, and let $\mathcal{L}$ be the real Lie algebra generated by constant $\mathfrak{g}$-valued functions and

$$z \longmapsto \mathrm{Ad}_z A = zAz^{-1}, \qquad A \in \mathfrak{g}.$$

These generators are infinitesimally available in $\mathcal{G}_+^{\mathcal{V}}(F)$ because the closure contains constants, z, and z^{-1}. Uniform Lie–Trotter and group commutator formulas imply that $e^{tY} \in \mathcal{G}_+^{\mathcal{V}}(F)$ for every $Y \in \mathcal{L}$ and $t \in \mathbb{R}$.

Fix a real basis $T_1, \ldots, T_{D^2-1}$ of $\mathfrak{g}$ and write

$$\mathrm{Ad}_z T_b = \sum_a q_{ab}(z) T_a. \tag{13}$$

The complexification of the adjoint representation of $\mathfrak{su}(D)$ is the irreducible adjoint representation of $\mathfrak{sl}_D(\mathbb{C})$. By Burnside's theorem, the complex associative algebra generated by the operators ad_A is the full endomorphism algebra; comparison of real dimensions gives $\mathrm{End}_{\mathbb{R}}(\mathfrak{g})$ before complexification. Applying its rank-one coordinate maps to Eq. (13), through finite linear combinations of nested commutators with constant functions, yields

$$q_{ab} T_c \in \mathcal{L} \qquad \text{for all } a, b, c.$$

Define the scalar-function algebra

$$\mathcal{A} = \{ f \in C(Z, \mathbb{R}) : fT \in \mathcal{L} \text{ for every } T \in \mathfrak{g} \}.$$

It contains 1 and all q_{ab}. It is closed under multiplication: for $f, g \in \mathcal{A}$ and any $T \in \mathfrak{g}$, write $T = \sum_j [A_j, B_j]$, which is possible because $[\mathfrak{g}, \mathfrak{g}] = \mathfrak{g}$, and observe that

$$fgT = \sum_j [fA_j, gB_j] \in \mathcal{L}.$$

The functions q_{ab} separate points of Z. Otherwise $\mathrm{Ad}_z = \mathrm{Ad}_w$ for some $z \neq w$, so z and w would differ by an element of the center of $\mathrm{SU}(D)$, contradicting Eq. (12). The real Stone–Weierstrass theorem gives

$$\overline{\mathcal{A}} = C(Z, \mathbb{R}), \qquad \overline{\mathcal{L}} = C(Z, \mathfrak{g}).$$

Continuity of the exponential and closedness of $\mathcal{G}_+^{\mathcal{V}}(F)$ imply that it contains e^Y for every $Y \in C(Z, \mathfrak{g})$.

Consequently, every null-homotopic continuous $V : Z \to \mathrm{SU}(D)$ belongs to $\mathcal{G}_+^{\mathcal{V}}(F)$. To see this, let $H : [0,1] \times Z \to \mathrm{SU}(D)$ be a homotopy from I_D to V. Uniform continuity provides a partition $0 = t_0 < \cdots < t_J = 1$ for which every $H(t_j, z) H(t_{j-1}, z)^{-1}$ lies in one logarithm neighborhood of I_D. Thus

$$H(t_j, z) H(t_{j-1}, z)^{-1} = e^{Y_j(z)}$$

with continuous $Y_j : Z \to \mathfrak{g}$, and $V = e^{Y_J} \cdots e^{Y_1} \in \mathcal{G}_+^{\mathcal{V}}(F)$.

Step 3: a sufficiently small centered cube satisfies these hypotheses. For $K_a = [-a, a]^3$, set

$$Z_a = E_n(K_a), \qquad \nu = (E_n)_{\#}\mu.$$

Because $E_n(\mathbf{0}) = I_D$ and E_n is continuous, there exists $a_{\log}(n) > 0$ such that, for $0 < a < a_{\log}(n)$, the principal logarithm is defined on Z_a and

$$\sup_{z \in Z_a} \|\log z\| < r_D.$$

We can shrink the cube once more to exclude distinct central multiples. Let

$$\Delta_D = \min_{\substack{\zeta^D = 1 \\ \zeta \neq 1}} |1 - \zeta| > 0.$$

Continuity gives $a_{\mathrm{ctr}}(n) > 0$ such that $\|z - I_D\| < \Delta_D/3$ for all $z \in Z_a$ whenever $0 < a < a_{\mathrm{ctr}}(n)$. If $z = \zeta w$ for $z, w \in Z_a$, then

$$|1 - \zeta| = \|z - w\| \leq \|z - I_D\| + \|w - I_D\| < \frac{2\Delta_D}{3},$$

which is impossible for nontrivial ζ. Hence Eq. (12) holds. Take

$$a_0(n) = \min\{a_{\log}(n), a_{\mathrm{ctr}}(n)\}.$$

It remains to show that ν has no atoms despite the Euler-coordinate degeneracy at $x_2 = 0$. For fixed $z \in Z_a$, define the real-analytic function

$$f_z(\boldsymbol{x}) = \|E_n(\boldsymbol{x}) - z\|_F^2.$$

This function is not identically zero because E_n is not constant on any nonempty open set. The zero set of a nonzero real-analytic function has three-dimensional Lebesgue measure zero. Since μ has a density,

$$\nu(\{z\}) = \mu\big(E_n^{-1}(\{z\})\big) = \mu\big(f_z^{-1}(0)\big) = 0.$$

Thus ν is nonatomic. In particular, the one-dimensional fibers caused by $e(x_1, 0, x_3) = R_z(x_1 + x_3)$ carry no probability mass.

Step 4: continuous maps on the signal image approximate any target law. Fix $\varepsilon > 0$. Compactness of $\mathcal{S}_n$ gives a finitely supported law

$$Q_M = \sum_{j=1}^{M} q_j \delta_{\rho_j}, \qquad W_2(Q, Q_M) < \varepsilon/3. \tag{14}$$

Because ν is nonatomic, there is a Borel partition $Z_a = A_1 \sqcup \cdots \sqcup A_M$ with $\nu(A_j) = q_j$. By inner regularity, choose pairwise disjoint compact sets $K_j \subset A_j$ such that

$$\nu\left(Z_a \setminus \bigcup_{j=1}^{M} K_j\right) < \delta.$$

The finite collection $K_1, \ldots, K_M$ has pairwise disjoint open neighborhoods O_j. By Urysohn's lemma, there are continuous $\eta_j : Z_a \to [0, 1]$ satisfying

$$\eta_j|_{K_j} = 1, \qquad \mathrm{supp}(\eta_j) \subset O_j.$$

Choose $U_j \in \mathrm{SU}(D)$ such that $U_j|0^n\rangle\langle 0^n|U_j^\dagger = \rho_j$, and choose a continuous path $\Gamma_j : [0, 1] \to \mathrm{SU}(D)$ from I_D to U_j. Define

$$V(z) = \prod_{j=1}^{M} \Gamma_j(\eta_j(z)).$$

The supports are disjoint, so at most one factor differs from the identity at any z. Therefore V is continuous and $V(z) = U_j$ on K_j. Moreover,

$$H(s, z) = \prod_{j=1}^{M} \Gamma_j(s\eta_j(z))$$

is a homotopy from I_D to V. Step 2 therefore places V in the positive-word closure.

Let Q_V be the law of $V(z)|0^n\rangle\langle 0^n|V(z)^\dagger$ for $z \sim \nu$. Couple this state to ρ_j whenever $z \in A_j$. The states agree on $\bigcup_j K_j$, and the trace-distance diameter of $\mathcal{S}_n$ is one, so

$$W_2(Q_V, Q_M) \le \sqrt{\delta}. \tag{15}$$

Choose $\delta < (\varepsilon/3)^2$. By Steps 1–2, a parameter-layer positive word $\widetilde{U}_{P,\Theta} : Z_a \to \mathrm{SU}(D)$ with finite P can satisfy

$$\sup_{z \in Z_a} \|\widetilde{U}_{P,\Theta}(z) - V(z)\| < \varepsilon/3.$$

Its pullback $U_{P,\Theta}(\boldsymbol{x}) = \widetilde{U}_{P,\Theta}(E_n(\boldsymbol{x}))$ has exactly the forward-only form in Eq. (2). For the reference state $\rho_0 = |0^n\rangle\langle 0^n|$,

$$d\Big(\widetilde{U}_{P,\Theta}(z)\rho_0\widetilde{U}_{P,\Theta}(z)^\dagger, V(z)\rho_0 V(z)^\dagger\Big) \le \|\widetilde{U}_{P,\Theta}(z) - V(z)\|.$$

Coupling through the same z therefore gives

$$W_2(Q_{P,\Theta}, Q_V) < \varepsilon/3.$$

Combining this inequality with Eqs. (14) and (15) proves $W_2(Q, Q_{P,\Theta}) < \varepsilon$. □

A.3 Why the Function-Level Argument Is Local

For completeness, we justify the locality limitation stated in the main text. When $n = 1$ and the signal variable ranges over all of $\mathrm{SU}(2) \simeq S^3$, a length-$P$ positive word

$$W_P(g) = A_P g A_{P-1} g \cdots A_1 g$$

has homotopy degree P: multiplication adds classes in $\pi_3(\mathrm{SU}(2)) \cong \mathbb{Z}$, while constant maps have class zero. The inversion map $g \mapsto g^{-1}$ has degree -1. Uniformly sufficiently close maps into $\mathrm{SU}(2)$ are homotopic, so positive words cannot converge uniformly to inversion on the complete group. Thus the inverse-synthesis and uniform function-density argument above cannot be extended unchanged to an unrestricted encoder image. This observation alone does not rule out every possible distribution-only universality statement on a larger domain.

B QDT and QGAN Loss Functions

QDT. The QDT baseline uses the fidelity-based loss of (Zhang et al., 2024):

$$\begin{aligned}
\mathcal{L}_{\mathrm{QDT}} &= 2r_{12} - r_{11} - r_{22}, \\
r_{12} &= 1 - \mathbb{E}_{x\sim P_g,\, y\sim P_d}\left[|\langle x|y\rangle|^2\right], \\
r_{11} &= 1 - \mathbb{E}_{x,x'\sim P_g}\left[|\langle x|x'\rangle|^2\right], \\
r_{22} &= 1 - \mathbb{E}_{y,y'\sim P_d}\left[|\langle y|y'\rangle|^2\right].
\end{aligned}$$

Here P_g and P_d denote the generated and target state distributions. The terms r_{12}, r_{11}, and r_{22} average the infidelity between distributions, within generated states, and within target states, respectively. This is the squared maximum mean discrepancy (MMD) (Gretton et al., 2012) with the fidelity kernel:

$$\begin{aligned}
\mathcal{L}_{\mathrm{QDT}} &= \mathbb{E}_{x,x'\sim P_g}[k(x,x')] + \mathbb{E}_{y,y'\sim P_d}[k(y,y')] - 2\mathbb{E}_{x\sim P_g,\, y\sim P_d}[k(x,y)], \\
k(x,y) &= |\langle x|y\rangle|^2.
\end{aligned}$$

QGAN. The discriminator assigns each input state a score $D \in [0,1]$, where $\langle Z_0\rangle_\psi$ is measured on its readout qubit after the discriminator circuit:

$$\begin{aligned}
D(|\psi\rangle) &= \frac{1 + \mathrm{Re}\langle Z_0\rangle_\psi}{2}, \\
\mathcal{L}_D &= \mathbb{E}_{\mathrm{fake}}[D(\psi)] - \mathbb{E}_{\mathrm{real}}[D(\psi)], \\
\mathcal{L}_G &= -\mathbb{E}_{\mathrm{fake}}[D(\psi)].
\end{aligned}$$

Training alternates between minimizing $\mathcal{L}_D$, which favors higher scores for target states than generated states, and minimizing $\mathcal{L}_G$, which raises the scores of generated states. The other network's parameters are held fixed during each update.

C Measurement Details for Correlated-Noise Estimation

The two-qubit dataset used in Section 4.4 starts from

$$|\Psi\rangle = \sqrt{0.2}\, e^{0.3\mathrm{i}}|00\rangle + \sqrt{0.5}\, e^{1.1\mathrm{i}}|01\rangle + \sqrt{0.3}\, e^{2.4\mathrm{i}}|11\rangle. \tag{16}$$

For each sample, the error type and angle are independently drawn according to

$$|\widetilde{\Psi}\rangle = \begin{cases} e^{-\mathrm{i}\delta X_1 X_2}|\Psi\rangle, & \text{with probability } p, \\ e^{-\mathrm{i}\delta Z_1 Z_2}|\Psi\rangle, & \text{with probability } 1-p, \end{cases} \qquad \delta \sim \mathrm{Unif}[-\delta_0, \delta_0],$$

where $X_1X_2 = X \otimes X$, $Z_1Z_2 = Z \otimes Z$, $p = 0.8$, and $\delta_0 = \pi/3$. Each sample is a pure state.

Let $c_1 = \sqrt{0.5}e^{1.1\mathrm{i}}$ denote the $|01\rangle$ coefficient of the initial state in Eq. (16). For the XX branch,

$$e^{-\mathrm{i}\delta X_1 X_2} = \cos\delta\, I - \mathrm{i}\sin\delta\, X_1X_2.$$

Because $X_1X_2|01\rangle = |10\rangle$ and the initial state contains no $|10\rangle$ component,

$$\langle 10|e^{-\mathrm{i}\delta X_1X_2}|\Psi\rangle = -\mathrm{i}c_1\sin\delta, \qquad P_{10}^{XX} = |c_1|^2\sin^2\delta.$$

The ZZ branch is diagonal in the computational basis and therefore cannot create a $|10\rangle$ component, so $P_{10}^{ZZ} = 0$ for the initial state in Eq. (16).

Averaging over the branch choice and the independently sampled angle gives

$$F_{10} := \mathbb{E}\Big[|\langle 10|\widetilde{\Psi}\rangle|^2\Big] = p|c_1|^2 m_\delta, \qquad m_\delta := \mathbb{E}_\delta[\sin^2\delta].$$

For $\delta \sim \mathrm{Unif}[-\delta_0, \delta_0]$,

$$m_\delta = \frac{1}{2\delta_0}\int_{-\delta_0}^{\delta_0} \sin^2\delta\, \mathrm{d}\delta = \frac{1}{2} - \frac{\sin(2\delta_0)}{4\delta_0}.$$

With $\delta_0 = \pi/3$ and $|c_1|^2 = 0.5$, this yields $m_\delta = 0.2932517$ and

$$F_{10} = 0.1466258p.$$

In particular, the physical value $p = 0.8$ corresponds to $F_{10}^{\mathrm{theory}} = 0.1173007$.

For N generated states $\{|\psi_j\rangle\}_{j=1}^N$, we estimate the mean marker population and then invert the calibration:

$$\widehat{F}_{10} = \frac{1}{N}\sum_{j=1}^{N}|\langle 10|\psi_j\rangle|^2, \qquad \widehat{p} = \frac{\widehat{F}_{10}}{|c_1|^2 m_\delta} = \frac{\widehat{F}_{10}}{0.1466258}.$$

Thus $\widehat{p}$ measures the effective XX mixing weight reproduced by the generated states.

D Experimental Details and More Experiments

D.1 Experimental Details

This section records the implementation and training settings for the six experiments in Section 4. Baseline loss formulas are collected in Appendix B, and the noise-estimation derivation is in Appendix C.

D.1.1 QuDDPM on the 0-Cluster Dataset

Data and circuits. The target dataset contains 100 training states around $|0\rangle$ and 1000 independent held-out states. Each training loss compares all 100 target states with 100 generated states. We evaluate $T = 1, \ldots, 20$ QuDDPM steps, using four Rot–C-R_x layers per step in our implementation. The baseline uses two auxiliary qubits with trainable rotations and selects the $|00\rangle$ measurement branch. The data qubit receives four trainable Rot gates per reverse step. Setting $P = 4T$ therefore matches number of trainable gates. Ancilla parameters and any learned latent scales are additional resources and are not included in this matching convention.

Training. QuDDPM samples Haar-random input states. Our generator samples classical noise from $[-0.1, 0.1]^3$ and uses three independent learnable scales. Both methods minimize the empirical Wasserstein objective in Eq. (5). For a fixed training set, each training run uses 2000 epochs with Adam (Kingma & Ba, 2015) at learning rate 0.005; circuit parameters are initialized uniformly in $[-0.1, 0.1]$.

Evaluation and resource counting. The Wasserstein evaluation compares 1000 generated states with the 1000 held-out target states. Figure 3 reports means, with error bars showing the observed minimum and maximum. If p_t denotes the conditional probability of the selected branch at reverse step t, QuDDPM's end-to-end success probability is $\prod_{t=1}^{T} p_t$. Our generator uses no post-selection, so its ideal success probability is one. A full loss for batches of M target and N generated states requires MN pairwise overlap estimates for our model and TMN across QuDDPM's denoising objectives. At $M = N = 100$, these counts are 10,000 and $10,000T$, respectively. They count pairwise estimates, not the number of measurement shots used for each estimate.

D.1.2 QDT AND QGAN ON THE REAL-AMPLITUDE CIRCLE

The target states in Eq. (8) have Bloch coordinates $(r_x, r_y, r_z) = (\sin 2\varphi, 0, \cos 2\varphi)$, so they form the x–z great circle. We use 100 independently sampled training states. QDT uses one data qubit and one ancilla with a 20-layer Rot–C-R_x generator. QGAN uses a 20-layer generator and a 16-layer discriminator of the same gate type. Our generator uses one qubit and $P = 20$ repetitions of data and trainable Rot gates.

All three methods use Adam at learning rate 0.005 for 2000 epochs, with parameters initialized uniformly in $[-0.1, 0.1]$. Both the QGAN generator and discriminator receive 2000 training epochs. QDT and QGAN take Haar-random pure states as inputs. The classical-noise training configuration uses $\boldsymbol{x} \sim \mathrm{Unif}([-0.1, 0.1]^3)$ and three independent learnable scales. Each experiment is repeated five times with different parameter initializations. The QDT and QGAN losses are specified in Appendix B.

D.1.3 ENTANGLEMENT-ENTROPY FITTING

We supervise the entropy of the first qubit's reduced density operator, $S(\rho_1) = -\operatorname{Tr}(\rho_1 \log_2 \rho_1)$, using

$$\mathcal{L}_S = \big(S(\rho_1(\boldsymbol{x}, \Theta)) - S_{\mathrm{tar}}\big)^2.$$

We test 11 target entropies $0, 0.1, \ldots, 1$ with 100 states per target. Our generator uses three-dimensional noise $\boldsymbol{x} \sim \mathrm{Unif}([-0.1, 0.1]^3)$ with three learnable scales, encoded as $\mathrm{Rot}(s_1 x_1, s_2 x_2, s_3 x_3)$, and $P = 20$ re-uploadings. The scales are optimized jointly with the circuit parameters. The QFM configurations use 20 layers. All four settings use Adam at learning rate 0.05 for up to 2000 epochs. We report the mean and standard deviation over three independent runs.

The QFM comparisons include the data-only circuit U_n, the ancilla-assisted circuit U_{n_a}, and a variant that switches between them across target steps. Our circuit form remains fixed throughout the target sequence. This experiment evaluates criterion-conditioned state generation by matching entropy, rather than recovering a specified distribution of complete states.

D.1.4 CORRELATED-NOISE PARAMETER ESTIMATION

Both methods use the same fixed set of 200 training states generated by the noise model in Appendix C. This finite training set gives an empirical estimate $\widehat{p} = 0.7998$. Evaluation uses $N \in \{1000, 2000, 5000\}$ generated states, with errors measured against the physical value $p = 0.8$.

Our generator uses $P = 20$ repetitions with C-R_x entangling gates, three-dimensional noise from $[-0.1, 0.1]^3$, and parameters initialized uniformly in $[-0.1, 0.1]$. LPQC uses 20 PQC layers with the same entangling gate and latent dimension, together with a two-layer, width-32 `tanh` flow network. Both methods use five independent training runs, 2000 epochs, and Adam at learning rate 0.005. The reported trainable parameter counts are 143 for our model and 5164 for LPQC.

The $|10\rangle$ populations are computed exactly from simulated statevectors; finite-shot and readout errors are not included. Thus the repeated-run errors characterize the generator's parameter recovery

under this protocol. We first compute $|\widehat{p}_r - p|$ in each run and then its mean and standard deviation across runs, rather than taking the absolute error of the mean estimate.

D.1.5 TEN-QUBIT QUANTUM PHASE LEARNING

The transverse-field Ising model uses the open-chain Hamiltonian

$$H(g) = -\sum_{j=1}^{n-1} Z_j Z_{j+1} - g \sum_{j=1}^{n} X_j,$$

where X_j and Z_j are Pauli operators on qubit j. We set $n = 10$ and draw 100 field strengths $g \sim \mathrm{Unif}[1.3, 1.5]$; their ground states form the training dataset. Our generator uses $P = 100$ re-uploadings, while LPQC uses 100 PQC layers. Our generator is trained with Eq. (5), three-dimensional noise $\boldsymbol{x} \sim \mathrm{Unif}([-0.1, 0.1]^3)$, and Adam at learning rate 0.05 for 1000 epochs. Each generator produces 100 new states for evaluation.

We evaluate the distribution of transverse magnetization $M_x = n^{-1} \sum_{j=1}^{n} \langle X_j \rangle$, reporting its mean and standard deviation together with trainable parameter counts.

D.1.6 CONDITIONAL GENERATION FROM DISJOINT LATENT INTERVALS

Tasks and interval labels. Section 4.6 tests conditional distribution selection using Eq. (6). A class is specified only by a disjoint interval I_c, from which all three latent coordinates are sampled independently. We train separate models for two circles and four finite-width clusters; they share neither circuit parameters nor latent scales.

For the circle model, the intervals and target Bloch vectors are

$$\begin{aligned} I_{\langle Z\rangle^2=0} &= [-0.3, -0.1] \cup [0.1, 0.3]: \quad \boldsymbol{r}_Z(\varphi) = (\cos\varphi, \sin\varphi, 0), \\ I_{\langle Y\rangle^2=0} &= [-0.3, -0.1] \cup [0.1, 0.3]: \quad \boldsymbol{r}_Y(\varphi) = (\sin\varphi, 0, \cos\varphi), \end{aligned} \qquad \varphi \sim \mathrm{Unif}[0, 2\pi).$$

For the cluster model, four intervals request clusters centered at

$$[-0.7, -0.5]\colon |0\rangle, \quad [-0.3, -0.1]\colon |1\rangle, \quad [0.1, 0.3]\colon |+\rangle, \quad [0.5, 0.7]\colon |-\rangle.$$

Each cluster is sampled around its center with angular scale $\sigma_\varphi = 0.005$ radians, truncated at $3\sigma_\varphi$; it is therefore a state distribution rather than repeated copies of its center state. Every condition contains 200 training states and 1000 independently sampled held-out target states.

Model and training. Both task-specific models use one qubit, a three-dimensional latent vector, and $P = 20$. We train for 2000 epochs with Adam at learning rate 0.005, drawing new latent samples from every interval at every epoch. If a task has K conditions, its objective is

$$\mathcal{L}_{\mathrm{cond}} = \frac{1}{K} \sum_{c=1}^{K} W_{2,c}^2,$$

where each $W_{2,c}^2$ is computed separately using the infidelity cost in Eq. (3). This prevents a model from matching only the pooled mixture while ignoring which interval requested which distribution.

E ABLATION STUDIES

We study four hyperparameters: the noise width w, the latent dimension d, the re-uploading number P, and the number L of trainable layers after each three-angle encoding (one encoding per repetition when $d = 3$). To remove the influence of learnable rescaling, we disable scale learning and fix every scale to one throughout training and evaluation. Noise coordinates are sampled independently as $x_k \sim \mathrm{Unif}[-w/2, w/2]$, so w is the full interval width. All four studies use one qubit. We use the 0-cluster for the width and dimension studies, the $\langle Y\rangle^2 = 0$ circle for the P study, and the $\langle Z\rangle^2 = 0$ circle for the L study, making the respective changes in the generated geometry easier to observe.

E.0.1 Noise Width

We fix $P = 20$, $L = 1$, and $d = 3$, and compare $w \in \{0.02, 0.2, 2, 2\pi\}$. Figure 8 shows that overly broad noise intervals lead to outputs spread far beyond the target 0-cluster. With the scales fixed, increasing w directly increases the input variance,

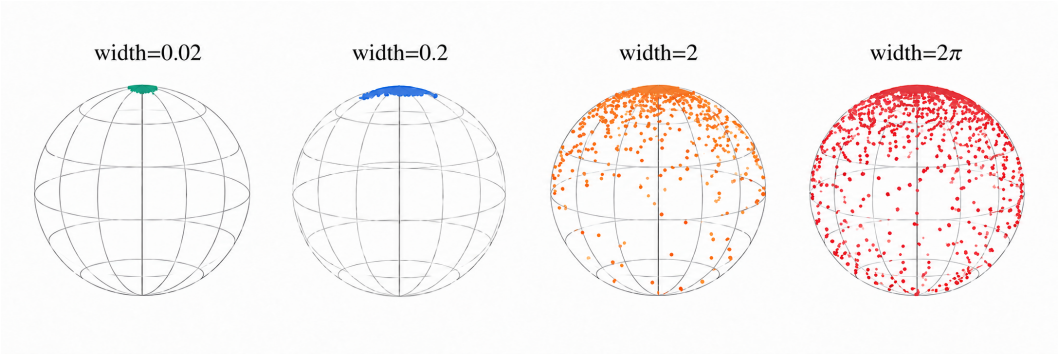


Figure 8: Qualitative noise-width comparison on the 0-cluster, with $P = 20$, $L = 1$, $d = 3$, and all scales fixed to one. Panels show $w = 0.02$, 0.2, 2, and 2π, respectively.

E.0.2 Noise Dimension at a Fixed Parameter Budget

We set $w = 0.5$ and keep 200 elementary encoding–parameter blocks, each containing a three-angle data Rot and one trainable Rot ($L = 1$). A d-dimensional vector is split into $d/3$ triples and encoded in successive blocks; P counts repetitions of the *complete* vector. We compare

$$(d, P) \in \{(3, 200), (30, 20), (300, 2), (600, 1)\}, \qquad P\frac{d}{3} = 200.$$

Thus every configuration has the same 200 trainable Rot gates and the same parameter count. Figure 9 shows poorer cluster recovery at larger d. This is a fixed-budget comparison of different dimensions and noise-reuse rates, rather than a dimension sweep at fixed P.

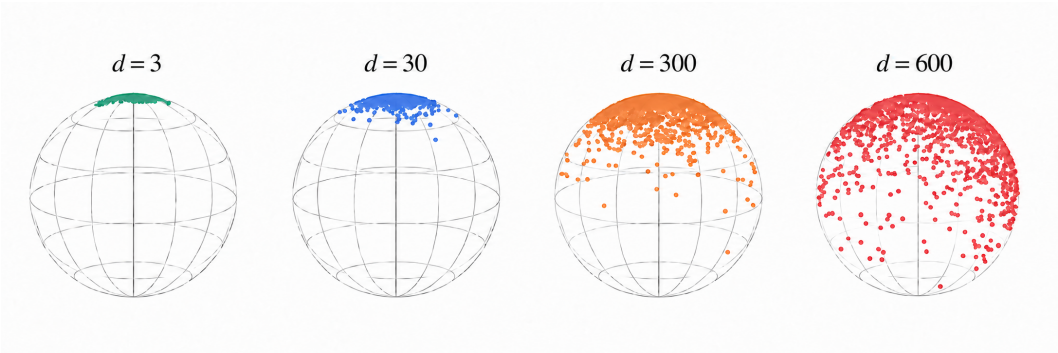


Figure 9: Qualitative noise-dimension comparison on the 0-cluster at $w = 0.5$. The four panels use $(d, P) = (3, 200)$, $(30, 20)$, $(300, 2)$, and $(600, 1)$, with the same total parameter count and all scales fixed to one.

E.0.3 Number of Re-Uploadings

We fix $w = 0.2$, $d = 3$, and $L = 1$, and compare $P \in \{1, 5, 20, 40\}$ on the $\langle Y \rangle^2 = 0$ circle. As shown in Fig. 10, larger P produces a progressively clearer target circle in these experiments. This agrees with the higher Fourier orders in Lemma 1 and the role of increasing P in Theorem 1. The theorem guarantees existence of sufficiently expressive circuits, not monotonic improvement of finite-budget training as P increases.

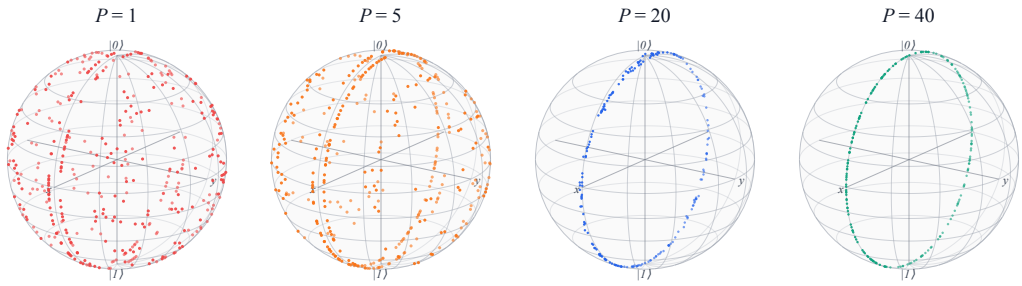


Figure 10: Qualitative re-uploading comparison on the $\langle Y \rangle^2 = 0$ circle. Panels use $P = 1$, 5, 20, and 40, with $w = 0.2$, $d = 3$, $L = 1$, and all scales fixed to one.

E.0.4 Parameter Layers within Each Repetition

We fix $w = 0.2$, $P = 20$, and $d = 3$, and compare $L \in \{0, 1, 2, 10\}$ on the $\langle Z \rangle^2 = 0$ circle. Figure 11 shows that $L = 0$ fails to recover the target circle, while $L = 1$ gives a close fit. Increasing L to 2 and 10 worsens the fit in the displayed results. This may reflect more difficult optimization of the deeper parameterization.

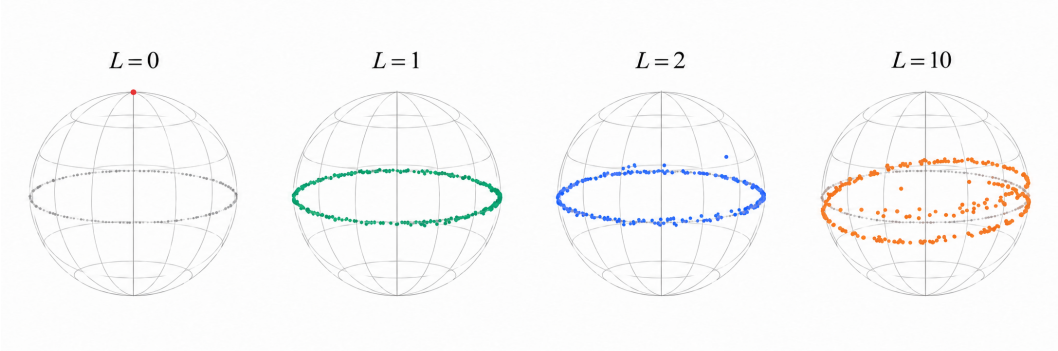


Figure 11: Qualitative parameter-layer comparison on the $\langle Z \rangle^2 = 0$ circle. Panels use $L = 0$, 1, 2, and 10, with $w = 0.2$, $P = 20$, $d = 3$, and all scales fixed to one.

Relation to prior analysis and our theorem. The deterioration at large width or dimension is qualitatively consistent with Wang et al. (2025): their analysis links larger input variance and deeper encodings of independent high-dimensional data in narrow circuits to loss of predictive information. Their setting concerns supervised prediction, whereas ours concerns distribution generation. The width trend also supports the practical relevance of the small-noise regime in Theorem 1.

F PURE AND MIXED QUANTUM STATE GENERATION

Architecture versus state purity. Our focus is *pure-quantum state generation*: the generator does not use a trainable classical neural network to map latent inputs to circuit parameters. Classical noise sampling, optimization, and the optimal-transport solver remain part of the workflow. This architectural distinction does not require the output states to be pure. The main-text implementation, experiments, and Theorem 1 concern distributions of pure states. Although averaging such states can produce a mixed density operator, this is not the same task as learning a distribution whose individual samples are mixed states: the average alone does not specify a state distribution.

Extension to mixed-state outputs. We can extend the generator by adding n_A ancillary qubits, applying the same re-uploading architecture on the enlarged register, and discarding the ancillas:

$$\rho_\Theta^S(\boldsymbol{x}) = \mathrm{Tr}_A\left[\widetilde{U}_\Theta(\boldsymbol{x})|0^{n+n_A}\rangle\langle 0^{n+n_A}|\widetilde{U}_\Theta(\boldsymbol{x})^\dagger\right].$$

An ancilla dimension of at least r suffices to represent a state of rank r; thus $n_A = n$ is sufficient for arbitrary n-qubit mixed states at the level of representation. Discarding the ancillas is unconditional, so generation still requires no post-selection and has ideal success probability one.

The expressivity argument also extends, provided the enlarged circuit satisfies the assumptions of Theorem 1, including its small-noise condition. Use a continuous choice of purifications of the target states and approximate the resulting pure-state distribution on the larger register. Partial trace cannot increase trace distance, and hence cannot increase the corresponding distributional W_2 distance. The reduced generated states therefore approximate the desired mixed-state distribution. This is an existence argument, not an efficient preparation or training guarantee; the extra ancillas and circuit depth also carry a resource cost.

Where distance estimation becomes harder. For mixed states, the classical optimal-transport problem in Eq. (5) is unchanged. The additional difficulty is obtaining its pairwise costs if we retain trace distance as the ground metric:

$$C_{ij} = d(\rho_i, \sigma_j)^2 = \frac{1}{4}\|\rho_i - \sigma_j\|_1^2.$$

Here W_2 compares distributions of density operators; it is not a separately defined quantum Wasserstein metric between two individual density operators. The SWAP test works for both pure and mixed states, but estimates their overlap through the ancilla outcome probability $p(0) = (1 + \mathrm{Tr}(\rho\sigma))/2$ (Ekert et al., 2002). For pure states this gives $d(\rho,\sigma)^2 = 1 - \mathrm{Tr}(\rho\sigma)$, as in Eq. (4). That identity does not hold for general mixed states: even $1 - \mathrm{Tr}(\rho^2)$ is nonzero when a mixed state is compared with itself. Mixed-state overlaps alone therefore do not determine the required trace-distance cost. Possible approaches are as follows.

1. **Tomography followed by classical computation.** Reconstruct each state and compute the eigenvalues of $\rho_i - \sigma_j$ to obtain its trace norm. For $D = 2^n$, an unrestricted density matrix has $D^2 - 1$ real parameters. Full-matrix storage uses $O(D^2)$ space per state, and a conventional dense eigendecomposition takes $O(D^3)$ operations per pair. Accurate tomography also requires many copies, with its cost depending on rank and precision (Haah et al., 2017). The generic approach is thus expensive in the number of qubits, although reconstructed states can be reused across pairs and additional structure can reduce the cost.
2. **Direct quantum trace-distance estimation.** Specialized algorithms can avoid reconstructing full density matrices, but their cost depends on rank, accuracy, and access to state preparation. For example, for states of rank at most r and additive error η, Wang & Zhang (2024) give algorithms using $\widetilde{O}(r/\eta^2)$ queries with purified-unitary access, or $\widetilde{O}(r^2/\eta^5)$ copies with sample access. These are different access models, and the former is stronger than merely receiving copies of a target state. Large rank or high accuracy can still make the cost substantial; low-rank problems may be much easier.
3. **Variational distance estimation.** One can use the characterization $d(\rho,\sigma) = \max_{0\le B\le I}\mathrm{Tr}[B(\rho-\sigma)]$ and optimize a parameterized measurement (Chen et al., 2022). This replaces full reconstruction with additional optimization and repeated measurements.

A restricted measurement family or incomplete optimization generally provides only a lower bound, rather than a certified exact cost. Repeating this procedure for many state pairs can be demanding, and efficient optimization is not guaranteed.

4. **A cheaper alternative ground metric.** If changing the loss is acceptable, the normalized Hilbert–Schmidt distance can be used instead:

$$d_{\mathrm{HS}}(\rho,\sigma)^2 = \frac{1}{2}\|\rho-\sigma\|_2^2 = \frac{1}{2}\operatorname{Tr}(\rho^2) + \frac{1}{2}\operatorname{Tr}(\sigma^2) - \operatorname{Tr}(\rho\sigma).$$

The purities and overlap are accessible by SWAP tests on independently prepared copies (Ekert et al., 2002). This agrees with squared trace distance on pure states, but generally differs on mixed states and therefore changes the Wasserstein objective. It shows that not every mixed-state loss necessarily has the cost of full tomography or trace-distance estimation.

A shared challenge for state-distribution generators. An empirical loss with M target and N generated states can require MN pairwise costs. Under the same state-access and accuracy requirements, obtaining these mixed-state costs is a shared problem for any generator trained or evaluated with the same trace-distance-based Wasserstein objective; it is not specific to data re-uploading. Other models can avoid explicitly computing this cost by using adversarial or observable-based losses, but then optimize a different criterion. Thus mixed-state distance estimation is a principal additional training bottleneck, not a restriction of the pure-quantum generator architecture to pure states.