# Physics-Constrained Quantum Circuit Born Machines for Synthetic Industrial Sensor Data Under Limited Training Samples: A Controlled Comparison With Classical Generators

[1]Pravat Satpathy, [2]Sonia Mohapatra, [3]Saswat Swain
[1]Researcher, [2]Researcher, [3]Ph.D Scholar
[1,2]Research and Development Division,
Tekstein Solutions Private Limited, Bhubaneswar, Odisha, India
[3]KSOM — KIIT Deemed to be University,
Bhubaneswar, Odisha, India
[1]analytics.pravat@gmail.com, [2]smctc98@gmail.com, [3]saswat.mck@gmail.com

*Abstract*—Industrial condition-monitoring programmes rarely possess the volume of labelled fault data that modern generative models require, and the few records of rare faults that do exist are often too sensitive to share. This paper develops and evaluates a physics-constrained quantum generative framework for producing synthetic bearing-condition records from very small training sets. A fourteen-qubit conditional quantum circuit Born machine (QCBM) learns the joint distribution of a fault label and four Gray-code-embedded vibration condition indicators (log-RMS, kurtosis, and the envelope-spectrum energy ratios at the outer-race and inner-race characteristic frequencies). Physical admissibility is enforced through a feasibility support derived from a parameter-envelope sweep of a rolling-element bearing impulse-response model, applied both as a soft penalty inside a maximum-mean-discrepancy objective and as a measurement post-selection at inference. The generator is trained with exact adjoint gradients and, in a hardware-realistic variant, with parameter-shift gradients estimated from 500 to 2000 measurement shots, and is sampled under depolarising gate noise. On training sets of 32 to 256 records with a 5% rare inner-race class, the framework is benchmarked against Gaussian-mixture, kernel-density, conditional-GAN and a parameter-matched classical energy model trained with the identical objective, using fidelity, diversity, physical validity, membership-inference privacy and train-on-synthetic/test-on-real utility. The projected QCBM eliminates physically impossible records, raises downstream macro-F1 by 0.24 on average over the unprojected circuit, lifts rare-class recall to the level of the best classical augmentation, and outperforms the parameter-matched classical energy model on fidelity at 32 to 64 records and on privacy at every size; it does not, however, surpass Gaussian-mixture or kernel-density generators on fidelity, coverage or utility on this four-indicator space, and all results are obtained by classical simulation of the circuit. The study therefore delimits where quantum generative models presently help, and identifies the feasibility constraint as the generator-agnostic ingredient that matters most under scarce data.

*Index Terms*—Quantum circuit Born machine, quantum generative models, synthetic sensor data, physics-constrained generation, bearing fault diagnosis, rare-fault augmentation, membership inference, maximum mean discrepancy, noisy intermediate-scale quantum.

## I. INTRODUCTION

Data-driven condition monitoring of rotating machinery is constrained in practice by a structural shortage of labelled fault records: bearings in critical drives are replaced before they fail, incipient faults are rare by design, and the vibration archives that exist are often confidential because they reveal operating schedules or supplier equipment. Classifiers trained on a few dozen windows generalise poorly, and the minority classes that matter most for maintenance planning, such as inner-race defects, have the fewest examples. Synthetic data generation is a natural remedy, but generative adversarial networks [27] and variational autoencoders [26] are themselves data-hungry, and simpler density estimators either impose strong parametric assumptions or memorise the training points.

Quantum generative models have been proposed for exactly this small-data regime. A quantum circuit Born machine (QCBM) represents a discrete distribution as the measurement statistics of a parameterised quantum state, and the normalisation and entanglement structure of that state have been argued to provide a useful inductive bias and, for some distributions, an expressivity that is hard to reproduce classically [3], [23]. Whether that bias helps on industrial sensor data, where samples are scarce but physically structured, has not been established with controlled experiments. Three gaps motivate this work: published QCBM studies use benchmark patterns or image digits [1], [7], [25] rather than sensor data with physical constraints; the physics of the measured process is ignored by the generator, so that a non-trivial fraction of synthetic records violate relations any real signal must obey; and comparisons with classical generators are rarely controlled

for parameter count and objective, so an observed difference cannot be attributed to the quantum parameterisation. This paper addresses these gaps for the generation of synthetic bearing-condition records from 32 to 256 training samples with a 5% rare inner-race class. The contributions are as follows.

1) A physics-constrained conditional QCBM framework in which the feasible support of a rolling-element bearing impulse-response model, obtained by sweeping the model over its admissible parameter envelope rather than from the training sample, is enforced as a soft penalty during training and as a measurement post-selection at inference. The post-selection is shown to be the closest feasible distribution in total variation (Proposition 1).

2) A Gray-code embedding of ordinal condition indicators onto qubits, together with a feature-space kernel for the maximum-mean-discrepancy objective, so that neighbouring quantisation levels differ in a single qubit and the loss respects the physical geometry of the indicators; the ablation shows the embedding to be harmless but of no measurable benefit at the depth studied.

3) A controlled benchmark against Gaussian-mixture, kernel-density and conditional-GAN generators and against a classical energy model with the same parameter count, objective, constraint and projection as the QCBM, evaluated on fidelity, diversity, physical validity, privacy leakage and downstream utility.

4) Hardware-realism studies covering parameter-shift training from finite measurement shots and depolarising gate noise at inference, with negative findings retained: on the four-indicator space considered, the quantum generator does not surpass classical density estimators, and its benefit is confined to comparisons against a matched classical energy model.

Section II reviews related work and Section III states the problem; Sections IV to VI present the framework, its formulation and the algorithm; Sections VII to IX give the setup, results and comparative analysis; Sections X and XI summarise the outcomes and conclude.

## II. RELATED WORK

Quantum generative modelling has developed along two lines. Explicit-amplitude models, principally the QCBM, treat the Born probabilities of a parameterised circuit as the model distribution and train them with sample-based divergences; Benedetti *et al.* introduced the data-driven circuit learning approach with a clipped negative log-likelihood and demonstrated it on bars-and-stripes patterns [1], Liu and Wang replaced the likelihood with a kernel maximum mean discrepancy (MMD) whose gradient admits a parameter-shift estimator from measurement samples [2], and Coyle *et al.* argued that Born machines inherit the classical hardness of sampling from certain circuits [3]. Adversarial formulations train a quantum generator against a classical or quantum discriminator [5], [6], and Zoufal *et al.* used a quantum GAN to load discretised distributions for option pricing [4]. Rudolph *et al.* generated handwritten digits with a QCBM prior embedded in a classical decoder on an ion-trap device [7]. Two recent assessments are directly relevant to the present study: Gili *et al.* examined whether Born machines generalise beyond the training set on cardinality-constrained distributions [25], and Hibat-Allah *et al.* proposed a framework that compares quantum against classical generative models on identical metrics and found no consistent quantum advantage on the datasets considered [23]. Rudolph *et al.* further showed that explicit-likelihood training of expressive circuits suffers from barren plateaus and that MMD-based implicit training is required at scale [28], consistent with the general trainability analysis of variational circuits [8], [9].

TABLE I SUMMARY OF CLOSELY RELATED WORK

| Work | Year | Generator and training | Data and evaluation | Limitation addressed here |
|---|---|---|---|---|
| Benedetti *et al.* [1] | 2019 | QCBM, clipped log-likelihood, data-driven circuit learning | Bars-and-stripes, Greenberger–Horne–Zeilinger; KL divergence | Benchmark patterns only; no physical constraint, no privacy or utility |
| Liu and Wang [2] | 2018 | QCBM, kernel MMD with parameter-shift gradients | Bars-and-stripes, Gaussian mixtures; MMD | No comparison controlled for capacity; no industrial data |
| Zoufal *et al.* [4] | 2019 | Quantum GAN, classical discriminator | Log-normal and bimodal 1-D distributions; KS statistic | One-dimensional targets; no constraint or privacy analysis |
| Rudolph *et al.* [7] | 2022 | QCBM prior + classical decoder on ion trap | Handwritten digits; FID | Hybrid pipeline hides circuit contribution; no sensor data |
| Gili *et al.* [25] | 2023 | QCBM, MMD; generalisation metrics | Cardinality-constrained bit strings | Constraint is combinatorial, not physical; no utility or privacy |
| Hibat-Allah *et al.* [23] | 2024 | QCBM vs. RNN, transformer, WGAN; MMD | Cardinality sets; quality–coverage metrics | No physics constraint; no downstream task or privacy |
| Zhang *et al.* [21] | 2020 | Classical deep GAN for imbalanced diagnosis | Bearing and gearbox spectra; accuracy | Thousands of samples; no validity, privacy or quantum baseline |
| This work | 2026 | Conditional QCBM, MMD + physics penalty, projection; matched classical control | Physics-based bearing indicators, N = 32–256; fidelity, diversity, validity, privacy, TSTR, shots, noise | — |

Classical synthetic sensor data for fault diagnosis has been generated with GANs on vibration spectra, typically to rebalance classes [21], and with SMOTE-type interpolation [20]; tabular generators such as CTGAN

handle mixed columns [18], and sample-level audit metrics were proposed by Alaa *et al.* [19]. Privacy leakage of generative models is quantified with membership-inference attacks, formulated for classifiers [16] and extended to GANs [17]. The envelope-analysis procedure and the impulse-response bearing model used here follow Randall and Antoni [14] and the Case Western Reserve University benchmark study of Smith and Randall [15], from which the fault-frequency ratios are taken; the kurtosis indicator follows Antoni [29].

Table I summarises the works most closely related to the proposed framework. None combines a physically derived feasibility constraint with a quantum generator, none evaluates privacy leakage of a quantum generator, and only [23] controls the classical comparison for objective and capacity. The present work fills this gap, and its evaluation design follows the recommendation of [23] that the same metrics be applied to quantum and classical models on identical data.

## III. PROBLEM STATEMENT

Let each monitored window of vibration be summarised by a fault label $c \in \{0, 1, 2\}$ (healthy, outer-race, inner-race) and a vector of $d = 4$ continuous condition indicators $\mathbf{f} = (f_1, \ldots, f_4)$. A plant possesses only $N$ labelled records $\{(c_i, \mathbf{f}_i)\}$, with $N$ between a few tens and a few hundred, drawn from an operating distribution $q$ with a strongly imbalanced class prior $\boldsymbol{\pi} = (0.60, 0.35, 0.05)$. The objective is to construct a generator from which an arbitrary number of synthetic records can be drawn such that (i) the synthetic joint distribution is close to $q$, (ii) the synthetic support covers the diversity of $q$ rather than a few modes, (iii) every synthetic record is physically admissible, (iv) the synthetic set does not reveal which records were in the training set, and (v) a fault classifier trained on synthetic records performs on real data as well as one trained on the real records, in particular for the rare class.

Four obstacles make this difficult: the sample size is below the regime in which neural generators are stable, so the generator must carry a strong inductive bias; the rare class may be represented by one to five records; physical relations among the indicators (an elevated outer-race envelope ratio cannot occur with Gaussian kurtosis) cannot be learnt from so few samples and must be supplied; and a quantum generator on a noisy intermediate-scale device [22] exposes its probabilities only through measurement shots perturbed by gate noise. Throughout, all quantum circuits are simulated classically at fourteen qubits and no claim about execution on quantum hardware is made.

## IV. PROPOSED FRAMEWORK

Fig. 1 shows the framework. A physics-parameterised bearing model supplies the records of the study (Section IV-A); a signal-processing chain extracts four indicators that are quantised and Gray-code embedded onto twelve qubits (Section IV-B); a parameter-envelope sweep of the same physical model defines the feasible support (Section IV-C); a conditional QCBM with two label qubits learns the joint distribution under a penalised MMD objective (Sections IV-D and V); and samples are drawn with feasibility projection and label conditioning (Section IV-E). The evaluation protocol is described in Section VII.

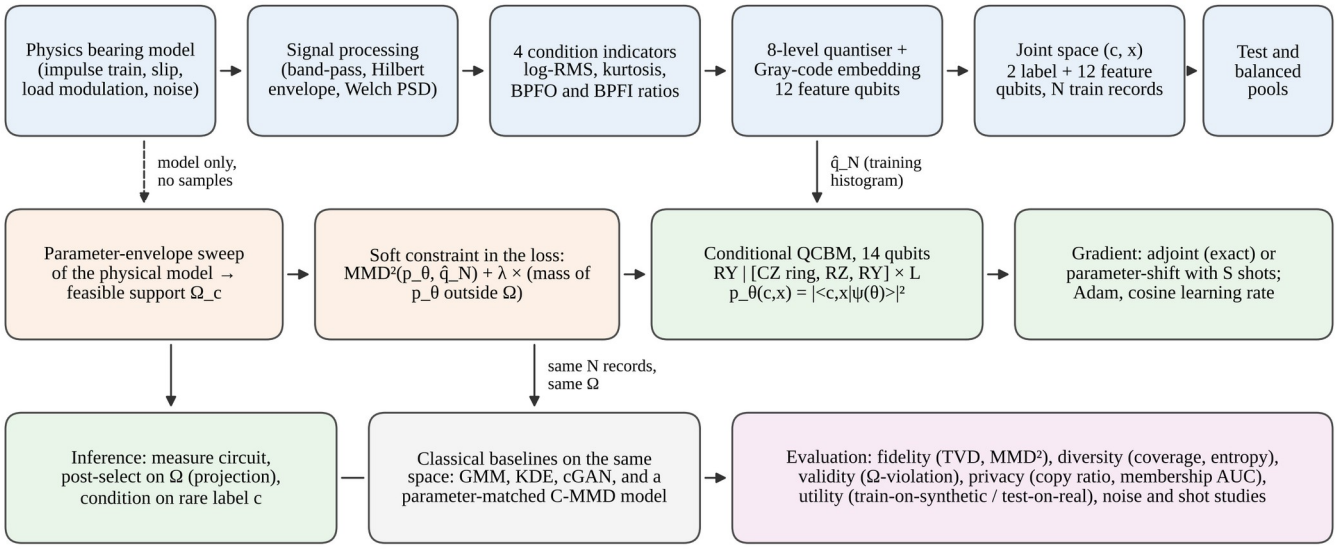



Fig. 1. Workflow of the physics-constrained quantum generative framework. The feasible support Ω is derived from the physical model alone (dashed arrow) and is shared by the quantum generator and by every classical baseline, so that the comparison isolates the generator.

### *A. Physics-Parameterised Vibration Source and Condition Indicators*

*Physical signal model.* Following the impulse-response description of Randall and Antoni [14], the acceleration measured on a bearing housing with a localised race defect is modelled as a train of structural resonance responses excited once per rolling-element passage,

$$x(t) = \Sigma_k A_k\, h(t - kT - \tau_k) + a_1 \sin(2\pi f_r t + \varphi_1) + a_2 \sin(4\pi f_r t + \varphi_2) + v(t) + n(t), \quad (1)$$

where $h(t) = \exp(-\zeta\omega_n t)\sin(\omega_n\sqrt{(1-\zeta^2)}\, t)$ for $t \geq 0$ is the response of the dominant housing resonance with natural frequency $\omega_n = 2\pi f_n$ and damping ratio $\zeta$; $T = 1/f_d$ is the nominal interval between impacts at the defect frequency $f_d$; $\tau_k \sim \mathcal{N}(0, (\sigma_s T)^2)$ is the cumulative random slip of the cage;

$A_k$ is the impact amplitude; $a_1$ and $a_2$ are shaft-harmonic amplitudes at the rotation frequency $f_r$; $v(t)$ is a sparse train of aperiodic interference impacts of random resonance; and $n(t)$ is white operating noise of standard deviation $\sigma_n$. For an outer-race defect the impacts occur in a fixed load zone, $A_k = A(1 + 0.1\varepsilon_k)$ with $\varepsilon_k \sim \mathcal{N}(0, 1)$, at $f_d$ = BPFO = 3.5848 $f_r$; for an inner-race defect the impact point rotates through the load zone, $A_k = A[1 + m\cos(2\pi f_r t_k + \varphi)]$, at $f_d$ = BPFI = 5.4152 $f_r$. The ratios correspond to the SKF 6205 drive-end bearing of the Case Western Reserve University test rig at $f_r$ = 29.95 Hz [15]. Healthy windows contain no periodic train. The severity $A/\sigma_n$, slip $\sigma_s$, resonance $f_n$, damping $\zeta$, modulation depth $m$, noise $\sigma_n$, shaft harmonics and a ±3% speed fluctuation that is unknown to the feature extractor are drawn per window from the deployment envelope of Table II.

$$\text{BPFO} = (N_b/2)\, f_r(1 - (d/D)\cos\alpha), \quad \text{BPFI} = (N_b/2)\, f_r(1 + (d/D)\cos\alpha), \tag{2}$$

with $N_b$ rolling elements of diameter $d$ on a pitch diameter $D$ at contact angle $\alpha$; for the SKF 6205 these evaluate to the ratios 3.5848 and 5.4152 used above [15].

*Condition indicators.* Each 4096-sample window at 12 kHz is reduced to four indicators by the standard envelope-analysis chain [14]: the base-10 logarithm of the root-mean-square value, the kurtosis, and the envelope-spectrum energy ratios

$$e_O = E_b(\text{BPFO})/E_{tot}, \quad e_I = E_b(\text{BPFI})/E_{tot}, \quad E_b(f_d) = \Sigma_{m=1}^{2} \int_{B_m} E(f)\, df, \tag{3}$$

$$\kappa = (1/n_w)\,\Sigma_i(x_i - \bar{x})^4 \,/\, [(1/n_w)\,\Sigma_i(x_i - \bar{x})^2]^2, \quad \mathbf{f} = (\log_{10}\text{RMS}, \kappa, \log_{10} e_O, \log_{10} e_I), \tag{4}$$

where $E(f)$ is the Welch power spectral density (2048-point Hann segments, 50% overlap) of the mean-removed Hilbert envelope of the signal band-pass filtered to 2.0 to 4.5 kHz around the structural resonance, $E_b$ sums the bands $B_m = [m f_d - \Delta,\ m f_d + \Delta]$ around the fundamental and second harmonic of the characteristic frequency $f_d$ with tolerance $\Delta$ = 8 Hz, and $E_{tot}$ is the total envelope-spectrum energy between 20 and 400 Hz. The band tolerance absorbs the speed fluctuation at the fundamental but not fully at the second harmonic, which is one of the sources of class overlap in Fig. 2(c).

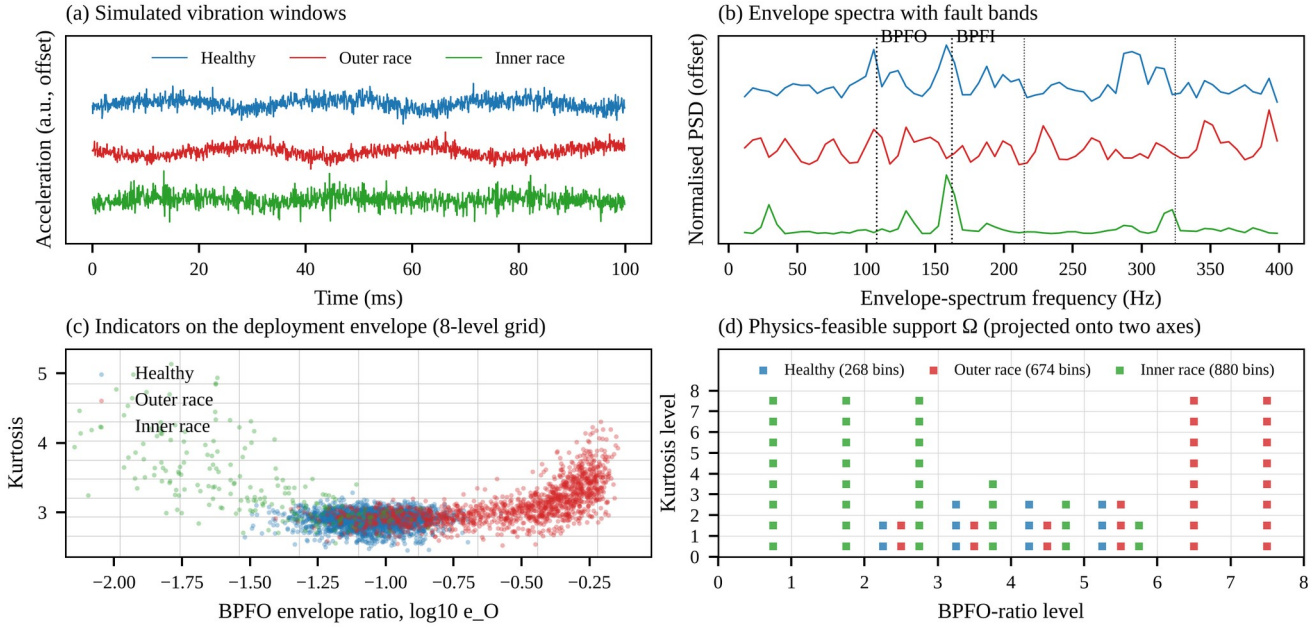



Fig. 2. Physics-parameterised data source. (a) Simulated windows of the three conditions. (b) Envelope spectra with the BPFO and BPFI bands (dotted lines mark the fundamentals and second harmonics). (c) Kurtosis against outer-race envelope ratio on the deployment envelope; grey lines are the fixed 8-level quantisation grid. (d) The physics-feasible support $\Omega_c$ derived from the sweep envelope, projected onto the same two axes.

### *B. Quantisation and Gray-Code Qubit Embedding*

Each indicator $f_j$ is mapped to an integer level $x_j \in \{0, \ldots, 7\}$ by a uniform quantiser between the 2nd and 98th percentiles of a balanced calibration pool drawn once from the deployment envelope; the grid is fixed for all experiments and generators. The joint record $(c, \mathbf{x})$ therefore lives on a discrete space of $3 \times 8^4 = 12\,288$ cells, encoded on $n$ = 14 qubits: two label qubits carrying $c$ in binary (the fourth label value is unphysical and is treated as infeasible) and three qubits per indicator. Ordinal levels are embedded with the reflected Gray code $g(x) = x \oplus \lfloor x/2 \rfloor$, so that adjacent levels differ in exactly one qubit; with plain binary encoding the transition from level 3 (011) to level 4 (100) flips all three qubits, which forces a nearest-neighbour entangling circuit to represent a smooth marginal with long-range correlations. The embedding is a fixed permutation of the computational basis and costs nothing on hardware; its measured effect is reported in Section VIII-E.

$$x_j = \min\{7, \max\{0, \lfloor 8(f_j - l_j)/(u_j - l_j)\rfloor\}\}, \quad g(x_j) = x_j \oplus \lfloor x_j/2 \rfloor, \quad j = 1, \ldots, 4, \tag{5}$$

### *C. Physics Feasibility Support*

The feasibility support $\Omega = \cup_c \{c\} \times \Omega_c$ is the set of (label, cell) pairs that the physical model of Section IV-A can produce anywhere in a deliberately widened parameter envelope (the sweep envelope of Table II). It is estimated once by drawing 60 000 balanced windows from the sweep envelope, quantising them on the fixed grid and retaining every cell observed at least twice, which yields $|\Omega_0|$ = 268, |

$\Omega_1| = 674$ and $|\Omega_2| = 880$ feasible cells, a union of 1435 of the 4096 feature cells. Two properties of this construction are essential. It is a property of the physical model and its admissible parameter ranges, not of the sampled training records, so it does not leak the training set; and it is strictly wider than the deployment distribution, only 1.8% of whose mass falls outside Ω because of the ±8 Hz band tolerance at the sweep extremes, so that the constraint removes impossible records without removing rare-but-real ones. Fig. 2(d) shows the support projected onto the kurtosis and outer-race-ratio axes. In a field deployment Ω would be supplied by a validated digital twin of the asset or by expert admissibility rules; the consequences of a misspecified Ω are discussed in Section VIII-G.

$$\Omega_c = \{\mathbf{x} \in \{0,\ldots,7\}^4 : n_{\text{sweep}}(c, \mathbf{x}) \geq 2\}, \quad \Omega = \cup_{c\in\{0,1,2\}} \{c\} \times \Omega_c, \tag{6}$$

### *D. Conditional Quantum Circuit Born Machine*

The generator is a hardware-efficient circuit acting on $n = 14$ qubits initialised in $|0\rangle^{\otimes n}$: a layer of single-qubit $R_Y(\theta)$ rotations followed by $L$ blocks, each consisting of a ring of controlled-Z gates between qubits $q$ and $(q + 1) \bmod n$, a layer of $R_Z$ rotations and a layer of $R_Y$ rotations. With $L = 4$ the circuit holds $P = n + 2nL = 126$ parameters and 56 CZ gates in 17 gate layers (the CZ gates of a ring commute and are scheduled as even edges, odd edges and the closing edge); transpiled by Qiskit 2.6 [24] onto a heavy-hexagon coupling map with the {RZ, SX, X, CZ} basis it requires 77 CZ gates and a depth of 66, the overhead coming from the SWAP gates that realise the ring on a lattice without a 14-cycle. Measuring all qubits in the computational basis yields the Born distribution (9). The two label qubits are not treated specially by the circuit; conditioning on a class is performed at inference by renormalising the measured joint distribution to that label, which on hardware is a post-selection of shots.

$$U(\theta) = \Pi_{l=1}^{L} \left[ \otimes_q R_Y(\theta_q^{(l,Y)}) \otimes_q R_Z(\theta_q^{(l,Z)}) \Pi_q \mathrm{CZ}_{q,q+1 \bmod n} \right] \otimes_q R_Y(\theta_q^{(0)}), \tag{7}$$

### *E. Inference: Projection and Rare-Label Conditioning*

At inference, measurement outcomes outside Ω are discarded and the remainder renormalised (feasibility projection), and rare-class records are obtained by additionally post-selecting the label qubits. The expected discard fraction equals the infeasible mass $\delta(\theta)$ of the trained circuit, which the soft penalty keeps small (7.1% on average over the sweep for the penalised circuit against 18.4% without the penalty). For the simulator, the exact Born vector is available and projection is a multiplication by the indicator of Ω followed by renormalisation; samples are then drawn from the projected vector.

$$p_\Omega(z) = \mathbf{1}_\Omega(z)\, p_\theta(z) / (1 - \delta), \quad \delta = \Sigma_{z\notin\Omega}\, p_\theta(z), \quad p_\Omega(\mathbf{x} \mid c) = p_\Omega(c, \mathbf{x}) / \Sigma_{\mathbf{x}'} p_\Omega(c, \mathbf{x}'). \tag{8}$$

## V. MATHEMATICAL FORMULATION

This section defines the objective, derives its gradients and states two propositions that are verified numerically in Section VIII.

*Born distribution and objective.* For parameters $\theta \in \mathbb{R}^P$ the circuit prepares $|\psi(\theta)\rangle = U(\theta)|0\rangle^{\otimes n}$ and the model distribution over joint indices $z = (c, \mathbf{x})$ is

$$p_\theta(z) = |\langle z|U(\theta)|0\rangle^{\otimes n}|^2, \quad z = (c, g(\mathbf{x})) \in \{0,1\}^{14}, \tag{9}$$

Let $\hat{q}_N$ be the empirical joint histogram of the $N$ training records. The training objective is the squared kernel MMD [10] between $p_\theta$ and $\hat{q}_N$ augmented with the physics penalty,

$$\mathrm{MMD}^2(p, q) = \mathbb{E}_{z,z'\sim p} k(z,z') - 2\mathbb{E}_{z\sim p,\, z'\sim q} k(z,z') + \mathbb{E}_{z,z'\sim q} k(z,z'), \tag{10}$$

$$\mathscr{L}(\theta) = \mathrm{MMD}^2(p_\theta, \hat{q}_N) + \lambda\, \Sigma_z\, \mathbf{1}_\Omega(z)\, p_\theta(z), \tag{11}$$

where $\mathbf{1}_\Omega(z)$ is the indicator of the infeasible set and $\lambda \geq 0$ is the penalty weight. The kernel is label-diagonal and acts on the integer feature coordinates rather than on the bit strings,

$$k((c,\mathbf{x}), (c',\mathbf{x}')) = \delta_{cc'}\, (1/|H|)\, \Sigma_{h\in H} \exp(-\|\mathbf{x} - \mathbf{x}'\|^2/2h^2), \quad H = \{0.5, 1, 2, 4\}, \tag{12}$$

with bandwidths $h \in \{0.5, 1, 2, 4\}$ grid units, so that the loss, like the Gray-code embedding, respects the ordinal geometry of the indicators. Writing $K$ for the $4096 \times 4096$ kernel matrix on feature cells and $P_c$, $Q_c$ for the class-$c$ blocks of $p_\theta$ and $\hat{q}_N$, the objective is the quadratic form

$$\mathscr{L}(\theta) = \Sigma_c\, [P_c^{\mathrm{T}} K \boldsymbol{P_c} - 2\boldsymbol{P_c}^{\mathbf{T}} \boldsymbol{K} Q_c + Q_c^{\mathrm{T}} K \boldsymbol{Q_c}] + \lambda\, 1_\Omega{}^{\mathbf{T}}\, \boldsymbol{p_\theta}, \tag{13}$$

and its gradient with respect to the probability vector is $g(z) = \partial\mathscr{L}/\partial p_\theta(z) = 2(K(P_c - Q_c))(\mathbf{x}) + \lambda \mathbf{1}_\Omega(z)$.

*Exact gradient by the adjoint method.* Because every term of (13) is a function of the probability vector, the chain rule gives $\partial\mathscr{L}/\partial\theta_k = \Sigma_z\, g(z)\, \partial p_\theta(z)/\partial\theta_k = \partial\langle\psi(\theta)|O|\psi(\theta)\rangle/\partial\theta_k$ with the diagonal observable $O = \Sigma_z\, g(z)|z\rangle\langle z|$. For a gate $U_k(\theta_k) = \exp(-i\theta_k G_k/2)$ with Pauli generator $G_k$, the adjoint recursion of Jones and Gacon [11] evaluates all $P$ derivatives from one forward and one backward sweep of the state,

$$\partial\mathscr{L}/\partial\theta_k = \mathrm{Im}\, \langle\lambda_k|\, G_k\, |\varphi_k\rangle, \quad k = 1, \ldots, P, \tag{14}$$

where $|\varphi_k\rangle = U_k \cdots U_1|0\rangle$ and $|\lambda_k\rangle = U_{k+1}^\dagger \cdots U_M^\dagger\, O|\psi\rangle$. The cost is $O(M2^n)$ per gradient, against $O(PM2^n)$ for parameter shift, which is what makes multi-seed sweeps at fourteen qubits affordable. The implementation was verified against PennyLane's backpropagation gradient to a maximum absolute error of 5.4e-18 and against central finite differences to 8.9e-13 (Section VII).

*Finite-shot gradient.* On hardware the probabilities are not accessible. For the MMD term the parameter-shift rule [13] gives the unbiased estimator of Liu and Wang [2],

$$\partial\mathscr{L}/\partial\theta_k \approx (1/S)\, \Sigma_{z\sim p_{\{\theta+\pi/2\, e_k\}}}\, \bar{k}(z) - (1/S)\, \Sigma_{z\sim p_{\{\theta-\pi/2\, e_k\}}}\, \bar{k}(z), \tag{15}$$

in which the witness $\bar{k}(z) = \mathbb{E}_{z'\sim p_{\{\theta\}}} k(z, z') - \mathbb{E}_{z'\sim \hat{q}_{\{N\}}} k(z, z') + \lambda \mathbf{1}_\Omega(z)/2$ is estimated from $S$ shots of the unshifted circuit and the two expectations from $S$ shots each of the $\pm\pi/2$-shifted circuits, for a total of $(2P + 1)S$ shots per step. The variance of each term decays as $O(1/S)$ because the witness is bounded by the kernel.

*Propositions.* Two properties of the physics constraint hold exactly and are verified numerically in Section VIII-C.

*Proposition 1 (projection is total-variation optimal).* Let $\delta = \Sigma_{z\notin\Omega}\, p(z)$ be the infeasible mass of a distribution $p$ and $p_\Omega(z) = \mathbf{1}_\Omega(z)p(z)/(1 - \delta)$ its feasibility projection. Then

$\mathrm{TV}(p, p_\Omega) = \delta$, and $\mathrm{TV}(p, r) \geq \delta$ for every distribution $r$ supported on $\Omega$.

*Proof.* $\mathrm{TV}(p, p_\Omega) = \tfrac{1}{2}[\delta + \Sigma_{z\in\Omega}\, p(z)(1/(1-\delta) - 1)] = \tfrac{1}{2}[\delta + \delta] = \delta$. For any $r$ on $\Omega$, $\mathrm{TV}(p, r) = \tfrac{1}{2}[\delta + \Sigma_{z\in\Omega}|p(z) - r(z)|] \geq \tfrac{1}{2}[\delta + |(1-\delta) - 1|] = \delta$ by the triangle inequality. ∎

*Proposition 2 (penalty bounds the infeasible mass).* At any parameter $\theta$ with objective value $\mathscr{L}(\theta)$, the infeasible mass satisfies $\delta(\theta) \leq \mathscr{L}(\theta)/\lambda$, because the MMD term is non-negative. Consequently a trained circuit with $\mathscr{L}(\theta^*) = \ell$ discards at most $\ell/\lambda$ of its shots under projection.

Proposition 1 justifies post-selection as the best possible feasibility repair for a given circuit, and Proposition 2 ties the hardware cost of that repair, the discard rate, to the final loss.

## VI. ALGORITHM AND COMPLEXITY

Algorithm 1 summarises training with the exact gradient; the finite-shot variant replaces line 6 by (15). Each adjoint gradient costs one forward and one backward statevector sweep over $M = P + nL$ gates on $2^n$ amplitudes, i.e. $O(M2^n) = O(nL2^n)$ time and $O(2^n)$ memory on the simulator; the kernel products cost $O(4 \times 4096^2)$ per step and dominate at $n = 14$. The parameter-shift variant costs $(2P + 1)S$ circuit executions per step with $O(S)$ classical post-processing. Sampling after training is $O(2^n)$ on the simulator and one shot per record on hardware, divided by $(1 - \delta)$ for the projection. Space complexity of the classical post-processing is $O(4096^2)$ for the kernel, which is independent of $N$.

**Algorithm 1: Physics-constrained conditional QCBM training and sampling**

**Input:** training records $\{(c_i, \mathbf{x}_i)\}_{i=1}^{N}$, feasibility mask $\mathbf{1}_\Omega$, kernel $K$, depth $L$, weight $\lambda$, steps $T$, learning rate $\eta_0$.

**Output:** parameters $\theta^*$; synthetic records.

1: Build $\hat{q}_N$; Gray-code permutation $\Pi$ of the basis; precompute $K\,Q_c$ for each class.

2: Initialise $\theta \sim \mathscr{U}[0, 2\pi)^P$; Adam [30] state $(m, v) \leftarrow 0$.

3: **for** $t = 0$ to $T$ **do**

4: $\quad \eta_t \leftarrow 0.005 + \tfrac{1}{2}\eta_0[1 + \cos(\pi t/T)]$.

5: $\quad |\psi\rangle \leftarrow U(\theta)|0\rangle$; $p \leftarrow \Pi^{-1}|\psi|^2$; $g \leftarrow 2K(P_c - Q_c) + \lambda \mathbf{1}_\Omega$.

6: $\quad \nabla_\theta \mathscr{L} \leftarrow$ ADJOINT($U$, $\theta$, $O = \mathrm{diag}(\Pi\, g)$) by (14).

7: $\quad \theta \leftarrow$ ADAM($\theta$, $\nabla_\theta \mathscr{L}$, $\eta_t$).

8: **end for**

9: $p_\Omega \leftarrow \mathbf{1}_\Omega \odot p \,/\, (1 - \delta)$; draw records from $p_\Omega$, or from $p_\Omega(\cdot|c)$ for a rare label $c$.

## VII. SIMULATION SETUP

*Software and hardware.* The physical model, feature extraction and classical baselines are implemented in Python 3.13 with NumPy 2.5, SciPy 1.18 and scikit-learn 1.9; the QCBM is a purpose-written NumPy statevector simulator with adjoint differentiation, Pauli-trajectory noise and the finite-shot estimator. PennyLane 0.45 [12] is used only as an independent cross-check of probabilities and gradients, Qiskit 2.6 [24] for a second statevector cross-check and hardware transpilation, and PyTorch 2.14 for the conditional GAN. A MATLAB script that re-implements the indicator chain and the distribution metrics is provided in the reproducibility bundle for cross-checking but was not executed by the authors, who had no MATLAB licence; all reported numbers come from the Python chain. The global seed is 20261026. All experiments ran on two x86-64 CPU cores with 7 GB of memory; one 300-step QCBM training run takes about 60 s with the adjoint gradient and about 13 min per 150 steps with the 2000-shot parameter-shift estimator.

*Data and splits.* A deployment pool of 30 000 windows drawn with the class prior (0.60, 0.35, 0.05) is split into a 24 000-record test set for all fidelity and privacy metrics, a 1000-record held-out set that supplies the non-members for the membership-inference attack, and a 5000-record pool from which five independent training subsets of each size $N \in \{32, 64, 128, 256\}$ are drawn. A balanced pool of 6000 windows is the test set for downstream classification. The expected number of inner-race records in a subset is $0.05N$; subsets without any inner-race record were retained because they represent the realistic case.

*Generators.* The proposed QCBM ($L = 4$, $\lambda = 0.02$, 300 Adam steps, cosine learning rate from 0.1 to 0.005) is evaluated with and without the soft penalty and with and without projection. The classical baselines are a class-conditional Gaussian mixture with one to three full-covariance components selected by BIC and a diagonal fallback for classes with fewer than four records (GMM), a class-conditional Gaussian kernel density with Scott's bandwidth (KDE), a conditional GAN with two 32-unit hidden layers in generator and discriminator trained for 3000 steps on standardised indicators (cGAN), and the parameter-matched classical energy model (C-MMD): $p_\varphi(z) \propto \exp[E_\varphi(b(z))]$ with a one-hidden-layer tanh network of width 8 on the ±1-encoded 14-bit string (129 parameters against 126 for the QCBM), trained with exactly the objective (11), the same $\lambda$, steps and learning-rate schedule, and projected in the same way. GMM, KDE and cGAN generate in the continuous indicator space and are quantised on the fixed grid; their projected variants ("+mask") use rejection sampling against $\Omega$. Every generator produces 20 000 synthetic records per run; the QCBM and C-MMD variants with $\lambda = 0$ and with projection only are also reported.

*Metrics.* Fidelity is the total variation distance (TVD) and the squared MMD of (10) between the synthetic histogram and the test histogram; the TVD sampling floor, obtained from two halves of the test set, is 0.138. Diversity is the coverage, the fraction of test cells with at least two occurrences that receive at least one synthetic record. Physical validity is the fraction of synthetic records outside $\Omega$. Privacy is the copy ratio, the fraction of synthetic records identical to a training record divided by the same fraction for test records (unity means the generator reproduces training cells no more often than the true distribution), and the ROC area of a black-box membership-inference attack in the style of LOGAN [17] that scores a record by its synthetic frequency and must separate the $N$ training records from $N$ held-out records (0.5 means no leakage). Utility is train-on-synthetic/test-on-real (TSTR) macro-F1 of a logistic-regression classifier and a

200-tree random forest fitted on 2000 synthetic records and tested on the balanced pool, with inner-race recall reported separately; train-on-real (TRTR) on the $N$ records is the reference. The hardware-realism studies use $S \in \{500, 2000\}$ shots per circuit execution for 150 steps on the $N = 64$, seed-0 subset, and a depolarising channel after every gate of the trained $N = 128$ circuit with two-qubit error probability $p_2 \in \{0, 10^{-3}, 3 \times 10^{-3}, 10^{-2}, 3 \times 10^{-2}\}$, $p_1 = p_2/10$, simulated as 400 Monte-Carlo Pauli trajectories of 50 shots each.

TABLE II MODEL AND SIMULATION PARAMETERS

| Parameter | Deployment envelope | Sweep envelope (for Ω) |
|---|---|---|
| Sampling rate / window | 12 kHz / 4096 samples | same |
| Shaft frequency $f_i$; speed fluctuation | 29.95 Hz; ±3% | 29.95 Hz; ±5% |
| Fault ratios BPFO, BPFI | 3.5848, 5.4152 | same |
| Impulse-to-noise ratio $A/\sigma_n$ | 0.5–4.0 | 0.3–8.0 |
| Noise $\sigma_n$; slip $\sigma_s$ | 0.7–1.3; 0.5–2% | 0.5–1.6; 0–3.5% |
| Resonance $f_n$; damping $\zeta$ | 2.6–3.4 kHz; 0.04–0.07 | 2.3–3.8 kHz; 0.03–0.09 |
| Modulation depth $m$ (inner race) | 0.5–0.9 | 0.3–1.0 |
| Interference impacts per window | 0–12 | 0–20 |
| Class prior (healthy, outer, inner) | 0.60, 0.35, 0.05 | balanced |
| Demodulation band; band tolerance Δ | 2.0–4.5 kHz; ±8 Hz | same |
| Quantiser | 8 levels, 2nd–98th percentile | fixed |
| Qubits; depth $L$; parameters $P$ | 14; 4; 126 | — |
| λ; steps; learning rate | 0.02; 300; 0.1 → 0.005 cosine | — |
| Kernel bandwidths $H$ | {0.5, 1, 2, 4} grid units | — |
| Synthetic records per run; seeds | 20 000; 5 | — |

## VIII. RESULTS AND DISCUSSION

### A. Training Behaviour and Cross-Check

Fig. 3(a) and (b) show the training trajectories on the $N = 64$, seed-0 subset. The exact-gradient QCBM reduces the MMD$^2$ to the training histogram from 0.156 to 0.025 in 300 steps; with $\lambda = 0$ the same circuit reaches 0.015, i.e. the penalty costs about 60% of the achievable fit while reducing the infeasible mass from 25.2% to 8.4%. The matched classical energy model reaches a lower training MMD$^2$ (0.005) because its softmax parameterisation can place mass on individual cells, but Section VIII-C shows that this comes with memorisation. The NumPy simulator agrees with PennyLane to 2.8e-17 in probability and 5.4e-18 in gradient, and with central finite differences to 8.9e-13, on a six-qubit, three-layer instance of the same ansatz, and the Qiskit statevector of the full fourteen-qubit trained circuit reproduces the NumPy probabilities exactly, so the adjoint implementation can be used as ground truth for the shot-based estimator.

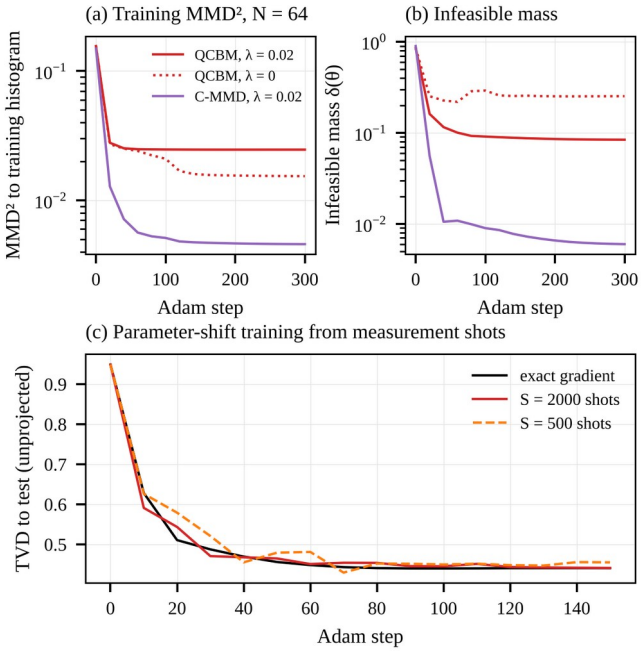



Fig. 3. Training behaviour on the $N = 64$, seed-0 subset. (a) MMD$^2$ to the training histogram for the penalised and unpenalised QCBM and the matched classical energy model. (b) Infeasible probability mass $\delta(\theta)$ during training. (c) Total variation distance to the test set during parameter-shift training from $S = 500$ and $S = 2000$ measurement shots per circuit execution, against the exact adjoint gradient with the same 150-step budget.

### B. Fidelity, Diversity and Physical Validity

Fig. 4 reports the five-seed means across the training-set sizes and Table III lists the full metric set at $N = 64$ and $N = 256$. Three observations follow. First, every unconstrained generator places a non-negligible share of its records on physically impossible cells: 19.1% for the unconstrained QCBM, 8.5% for the KDE, 8.2% for the cGAN and 4.2% for the GMM at $N = 64$ (Fig. 4(c)); the soft penalty reduces the QCBM figure to 7.1% and the projection removes it for every generator. Second, the projection improves rather than merely sanitises the quantum generator: the TVD of the QCBM falls from 0.470 to 0.428 at $N = 64$ and from 0.383 to 0.341 at $N = 256$, which is the mechanism of Proposition 1 at work, since the discarded mass (20.0% and 17.0%) was entirely outside the support of the true distribution. Third, the projected QCBM does not approach the Gaussian-mixture or kernel-density baselines on this indicator space. At $N = 256$ the projected QCBM reaches TVD 0.341 and coverage 0.61 against 0.211 and 0.80 for the masked GMM, whose TVD is 0.07 above the sampling floor of 0.138. The class-conditional indicator distributions of the physical model are close to Gaussian (Fig. 2(c)), so a mixture of Gaussians carries the correct inductive bias that the QCBM and the C-MMD model must discover through a kernel objective on a discretised grid.

Against the parameter-matched classical energy model the picture is different. With the same objective, penalty, projection and optimiser, the QCBM attains a lower TVD at

$N$ = 32 (0.521 against 0.700) and $N$ = 64 (0.472 against 0.540), higher coverage (0.48 against 0.34 at $N$ = 32), the two are equal at $N$ = 128 (0.420 against 0.419), and the classical model is lower at $N$ = 256 (0.364 against 0.406). The quantum parameterisation therefore behaves as a smoother, better-regularised model of the same size only in the scarcest regime, which is the regime the study targets; the advantage is relative to a classical model of the same family, disappears with more records, and is not an advantage over the best classical density estimator.

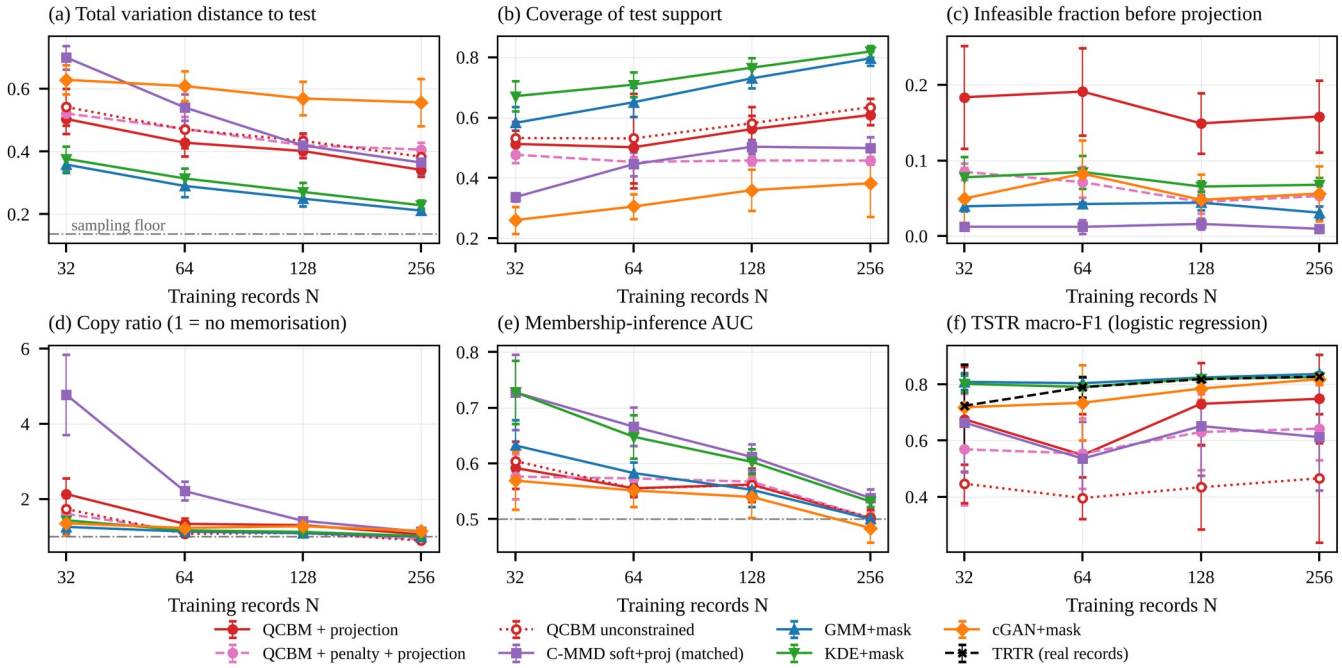



Fig. 4. Training-set-size sweep (mean ± standard deviation over five independent training subsets per $N$). (a) Total variation distance to the 24 000-record test histogram; the dash-dotted line is the two-half sampling floor. (b) Coverage of the test support. (c) Fraction of records on physically infeasible cells before projection or masking. (d) Copy ratio relative to the natural coincidence rate. (e) Membership-inference AUC. (f) Train-on-synthetic/test-on-real macro-F1 of a logistic-regression classifier, with train-on-real as reference.

TABLE III FULL METRIC SET AT N = 64 (TOP) AND N = 256 (BOTTOM), MEAN ± STANDARD DEVIATION OVER SUBSETS

| Generator | TVD ↓ | Coverage ↑ | Infeasible ↓ | Copy ratio | MI AUC | LR F1 ↑ | RF F1 ↑ | Inner recall ↑ |
|---|---|---|---|---|---|---|---|---|
| **N = 64** | | | | | | | | |
| Real records only (TRTR) | – | – | – | – | – | 0.79 ± 0.04 | 0.58 ± 0.16 | 0.66 ± 0.10 |
| **QCBM, penalty + projection** | 0.472 ± 0.038 | 0.453 ± 0.017 | 0.000 | 1.14 ± 0.09 | 0.57 ± 0.01 | 0.55 ± 0.12 | 0.39 ± 0.02 | 0.35 ± 0.28 |
| QCBM, penalty only | 0.488 ± 0.037 | 0.470 ± 0.025 | 0.071 | 1.06 ± 0.07 | 0.57 ± 0.01 | 0.32 ± 0.07 | 0.31 ± 0.04 | 0.02 ± 0.05 |
| QCBM, projection only | 0.428 ± 0.043 | 0.502 ± 0.135 | 0.000 | 1.34 ± 0.14 | 0.56 ± 0.01 | 0.55 ± 0.15 | 0.60 ± 0.18 | 0.16 ± 0.33 |
| QCBM, unconstrained | 0.470 ± 0.060 | 0.531 ± 0.148 | 0.191 | 1.08 ± 0.07 | 0.56 ± 0.02 | 0.40 ± 0.07 | 0.39 ± 0.06 | 0.00 ± 0.00 |
| C-MMD, penalty + projection | 0.540 ± 0.042 | 0.445 ± 0.039 | 0.000 | 2.21 ± 0.25 | 0.67 ± 0.03 | 0.54 ± 0.13 | 0.48 ± 0.04 | 0.12 ± 0.27 |
| GMM + mask | 0.290 ± 0.035 | 0.651 ± 0.049 | 0.000 | 1.14 ± 0.13 | 0.58 ± 0.02 | 0.80 ± 0.02 | 0.80 ± 0.04 | 0.69 ± 0.06 |
| GMM | 0.297 ± 0.035 | 0.695 ± 0.055 | 0.042 | 1.09 ± 0.12 | 0.59 ± 0.02 | 0.79 ± 0.04 | 0.80 ± 0.03 | 0.65 ± 0.09 |
| KDE + mask | 0.314 ± 0.032 | 0.709 ± 0.041 | 0.000 | 1.17 ± 0.09 | 0.65 ± 0.04 | 0.79 ± 0.04 | 0.81 ± 0.01 | 0.66 ± 0.09 |
| KDE | 0.337 ± 0.033 | 0.759 ± 0.044 | 0.085 | 1.08 ± 0.08 | 0.65 ± 0.04 | 0.79 ± 0.03 | 0.76 ± 0.10 | 0.65 ± 0.09 |
| cGAN + mask | 0.608 ± 0.048 | 0.305 ± 0.041 | 0.000 | 1.23 ± 0.21 | 0.55 ± 0.03 | 0.73 ± 0.13 | 0.73 ± 0.13 | 0.65 ± 0.31 |
| cGAN | 0.614 ± 0.050 | 0.321 ± 0.049 | 0.082 | 1.13 ± 0.18 | 0.55 ± 0.02 | 0.74 ± 0.11 | 0.70 ± 0.14 | 0.65 ± 0.26 |
| **N = 256** | | | | | | | | |
| Real records only (TRTR) | – | – | – | – | – | 0.83 ± 0.01 | 0.82 ± 0.02 | 0.76 ± 0.03 |
| **QCBM, penalty + projection** | 0.406 ± 0.023 | 0.457 ± 0.013 | 0.000 | 0.93 ± 0.07 | 0.50 ± 0.01 | 0.64 ± 0.11 | 0.43 ± 0.04 | 0.48 ± 0.27 |
| QCBM, penalty only | 0.418 ± 0.023 | 0.476 ± 0.014 | 0.053 | 0.88 ± 0.07 | 0.50 ± 0.01 | 0.50 ± 0.17 | 0.37 ± 0.08 | 0.38 ± 0.36 |
| QCBM, projection only | 0.341 ± 0.021 | 0.609 ± 0.032 | 0.000 | 1.06 ± 0.04 | 0.50 ± 0.02 | 0.75 ± 0.16 | 0.59 ± 0.20 | 0.60 ± 0.33 |
| QCBM, unconstrained | 0.383 ± 0.030 | 0.635 ± 0.029 | 0.158 | 0.90 ± 0.05 | 0.50 ± 0.01 | 0.47 ± 0.23 | 0.40 ± 0.09 | 0.18 ± 0.31 |
| C-MMD, penalty + projection | 0.364 ± 0.038 | 0.499 ± 0.037 | 0.000 | 1.13 ± 0.04 | 0.54 ± 0.02 | 0.61 ± 0.19 | 0.52 ± 0.13 | 0.30 ± 0.40 |
| GMM + mask | 0.211 ± 0.007 | 0.796 ± 0.024 | 0.000 | 1.00 ± 0.03 | 0.50 ± 0.01 | 0.84 ± 0.01 | 0.84 ± 0.01 | 0.79 ± 0.02 |
| GMM | 0.215 ± 0.009 | 0.853 ± 0.028 | 0.031 | 0.97 ± 0.03 | 0.50 ± 0.01 | 0.83 ± 0.01 | 0.82 ± 0.02 | 0.77 ± 0.04 |
| KDE + mask | 0.228 ± 0.016 | 0.819 ± 0.020 | 0.000 | 1.02 ± 0.03 | 0.53 ± 0.01 | 0.83 ± 0.01 | 0.82 ± 0.01 | 0.75 ± 0.03 |
| KDE | 0.249 ± 0.012 | 0.876 ± 0.016 | 0.068 | 0.95 ± 0.03 | 0.53 ± 0.01 | 0.82 ± 0.01 | 0.82 ± 0.00 | 0.75 ± 0.03 |
| cGAN + mask | 0.556 ± 0.075 | 0.383 ± 0.111 | 0.000 | 1.15 ± 0.08 | 0.48 ± 0.03 | 0.82 ± 0.02 | 0.81 ± 0.01 | 0.77 ± 0.05 |
| cGAN | 0.561 ± 0.078 | 0.398 ± 0.122 | 0.056 | 1.09 ± 0.08 | 0.49 ± 0.03 | 0.81 ± 0.03 | 0.80 ± 0.03 | 0.74 ± 0.08 |

### *C. Privacy Leakage and Verification of the Propositions*

The copy ratio (Fig. 4(d)) separates memorising from generalising generators. The matched classical energy model reproduces training cells 4.8 times more often than the true distribution does at $N = 32$ and 2.2 times at $N = 64$, with a membership-inference AUC of 0.73 and 0.67; the kernel density, which is centred on the training points by construction, has the highest attack AUC among the classical generators (0.74 at $N = 32$) although its copy ratio stays at 1.33. The projected QCBM stays at a copy ratio of 1.14 and an AUC of 0.57 at $N = 64$, comparable with the masked GMM (1.14, 0.58) and the masked cGAN (1.23,

0.55). Projection without the penalty raises the QCBM copy ratio to 2.13 at $N = 32$, because renormalisation concentrates mass on the feasible cells that the training records occupy; the soft penalty keeps the mass spread. At $N = 256$ all attacks are close to chance (AUC 0.48 to 0.54 across generators).

Proposition 1 was verified on the trained circuits: for the $N = 64$, seed-0 circuit the measured $\mathrm{TV}(p_\theta, p_\Omega)$ is 0.0839 against an infeasible mass $\delta = 0.0839$, equal to machine precision, while the TV to the uniform distribution on $\Omega$, another feasible candidate, is 0.764. Proposition 2 holds but is loose at the small penalty weight used: the final loss $\ell = 0.0263$ gives the bound $\delta \le \ell/\lambda = 1.32$, whereas the actual $\delta$ is 0.084; at $\lambda = 0.5$ the ablation of Section VIII-E gives $\ell/\lambda = 0.065$ against $\delta = 0.000$, so the bound becomes informative exactly when shots are expensive and a large $\lambda$ is chosen.

### *D. Downstream Utility and Rare-Fault Augmentation*

Fig. 4(f) shows train-on-synthetic/test-on-real macro-F1. Synthetic records from the masked GMM and KDE match or exceed training on the real records at every $N$ (0.80 against TRTR 0.79 at $N = 64$), and the masked cGAN comes within 0.01 of it at $N = 256$. The QCBM is the weakest generator on this criterion: the projected circuit yields 0.55 at $N = 64$ and 0.73 at $N = 128$, and the unconstrained circuit only 0.40, because its records off the physical support act as label noise for the classifier. The projection is therefore worth 0.15 to 0.30 in macro-F1, and the projection applied to the penalised circuit 0.17 on average, but neither closes the gap to the classical density estimators. The random-forest results in Table III show the same ordering between generator families, with larger seed-to-seed variation at $N \le 64$, where the forest overfits the discretised indicators.

Fig. 5(a) isolates the rare class. The downstream forest is trained on the $N$ real records plus enough synthetic inner-race records to equal the healthy count, with one to four real inner-race records available at $N = 64$ and three to eight at $N = 128$. Conditioning the projected QCBM on the inner-race label yields an inner-race recall of $0.76 \pm 0.15$ at $N = 64$ and $0.90 \pm 0.01$ at $N = 128$, compared with $0.22 \pm 0.30$ and $0.72 \pm 0.06$ without augmentation, $0.72 \pm 0.08$ and $0.74 \pm 0.09$ for SMOTE, $0.78 \pm 0.03$ and $0.82 \pm 0.03$ for the class-conditional GMM, and $0.84 \pm 0.02$ and $0.83 \pm 0.02$ for an oracle that adds real inner-race records. The quantum generator therefore lifts the rare-class recall to the level of SMOTE and the conditional GMM at $N = 64$ and above the oracle at $N = 128$, but its synthetic inner-race records are broader than the true class: the inner-race precision is $0.65 \pm 0.12$ and $0.78 \pm 0.06$ against $0.93 \pm 0.04$ and $0.90 \pm 0.05$ for the GMM, so that its macro-F1 ($0.72 \pm 0.08$ and $0.80 \pm 0.02$) remains below the GMM ($0.81 \pm 0.01$ and $0.81 \pm 0.01$). For a maintenance policy that weights a missed inner-race fault far above a false alarm this trade is acceptable; for a balanced cost it is not.

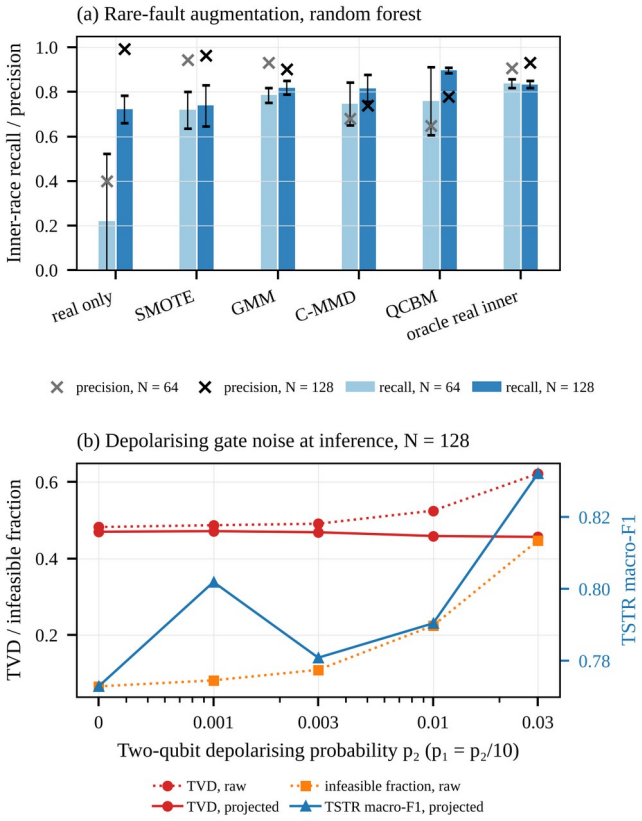



Fig. 5. (a) Inner-race recall of a random forest trained on the $N$ real records plus synthetic inner-race records from each augmentation method (mean ± standard deviation over five subsets); the oracle adds real records and bounds what augmentation can achieve. (b) Depolarising gate noise at inference for the $N = 128$ circuit: TVD before and after projection, infeasible fraction before projection, and projected TSTR macro-F1 against the two-qubit error probability.

### *E. Ablation: Depth, Penalty Weight, Encoding and Initialisation*

Table IV reports the ablation on the $N = 64$ subsets; the depth, penalty, initialisation and encoding variants use three seeds and the two reference rows the five seeds of the sweep, and seed-to-seed standard deviations of the projected TVD are about 0.05, so differences below that level are not resolved. Depth behaves as a capacity control rather than as a fidelity lever: the training $\mathrm{MMD}^2$ falls monotonically from 0.0282 at $L = 1$ (42 parameters) to 0.0195 at $L = 6$ (182 parameters), but the projected TVD to the test distribution does not improve (0.421 at $L = 1$ against 0.494 at $L = 6$) and the copy ratio rises from 0.99 to 1.37; the additional expressivity is spent on memorising the 64 records, in line with the generalisation concerns raised in [25]. Raising $\lambda$ from 0.02 to 0.1 and 0.5 lowers the infeasible mass before projection from 0.071 to 0.010 and 0.000, but the coverage collapses from 0.453 to 0.359 and 0.224 and the macro-F1 from 0.55 to 0.26: a strong penalty drives the optimiser to a few safe cells, so $\lambda$ should be chosen from the cost of a discarded shot on the target hardware and otherwise kept small or zero. The marginal-matching initialisation (TVD 0.453) and the binary encoding (0.468) are within seed variation of the default; the Gray code therefore has no measurable benefit at this depth and sample size, a negative result retained here. A sharper kernel with bandwidths {0.25, 0.5, 1, 2} leaves the

projected TVD of the QCBM unchanged (0.467) and worsens that of the C-MMD model from 0.540 to 0.605 with a copy ratio of 2.66; the kernel width moves the classical energy model along the fidelity–memorisation trade-off but does not close the gap to the density estimators for either model.

TABLE IV ABLATION ON THE N = 64 SUBSETS (MEAN OVER THREE SEEDS, FIVE FOR THE TWO REFERENCE ROWS; TVD, COVERAGE AND F1 AFTER PROJECTION; INFEASIBLE MASS BEFORE PROJECTION)

| Configuration | MMD² train | TVD | Coverage | Infeasible (raw) | Copy ratio | LR F1 |
|---|---|---|---|---|---|---|
| L = 4 (126 par.), λ = 0, five seeds | 0.0148 | 0.428 | 0.502 | 0.191 | 1.34 | 0.55 |
| L = 1 (42 par.) | 0.0282 | 0.421 | 0.499 | 0.090 | 0.99 | 0.45 |
| L = 3 (98 par.) | 0.0263 | 0.431 | 0.481 | 0.090 | 1.04 | 0.57 |
| L = 4 (126 par.), λ = 0.02, five seeds | 0.0230 | 0.472 | 0.453 | 0.071 | 1.14 | 0.55 |
| L = 6 (182 par.) | 0.0195 | 0.494 | 0.491 | 0.080 | 1.37 | 0.46 |
| L = 4, λ = 0.1 | 0.0290 | 0.536 | 0.359 | 0.010 | 1.11 | 0.34 |
| L = 4, λ = 0.5 | 0.0323 | 0.539 | 0.224 | 0.000 | 1.03 | 0.26 |
| L = 4, marginal-matching init. | 0.0259 | 0.453 | 0.475 | 0.083 | 1.11 | 0.61 |
| L = 4, binary encoding | 0.0253 | 0.468 | 0.466 | 0.094 | 1.10 | 0.59 |

### *F. Finite Measurement Shots and Gate Noise*

Fig. 3(c) shows parameter-shift training from measurement samples. With $S$ = 2000 shots per circuit execution, i.e. 506,000 shots per step, the circuit reaches a TVD of 0.441 after 150 steps, indistinguishable from the 0.441 of the exact adjoint gradient with the same step budget, and with $S$ = 500 the TVD is 0.455; after projection the three circuits give TVDs of 0.425, 0.429 and 0.439. The estimator is unbiased and its variance is already small enough at 2000 shots that the stochastic gradient does not slow convergence, so nothing in the training procedure depends on access to exact probabilities. The practical implication is that one training run costs of the order of 76 million shots, within reach of current cloud devices but far from negligible against the milliseconds that the classical baselines need.

Fig. 5(b) shows the effect of depolarising gate noise on the trained $N$ = 128 circuit (56 CZ and 126 single-qubit gates). Without projection the noise is destructive: at $p_2 = 10^{-2}$ the infeasible fraction rises from 6.6% to 22.5% and the TSTR macro-F1 falls from 0.59 to 0.25, and at $p_2 = 3 \times 10^{-2}$ almost half of the records (44.7%) are physically impossible. With projection the TVD varies by only 0.015 over the whole range and the projected macro-F1 rises from 0.77 to 0.83, with coverage increasing from 0.46 to 0.64. Depolarising noise spreads the Born distribution, and the part of that spread that stays on the physical support behaves as a smoothing of a circuit that is otherwise too concentrated; this observation comes from a single trained circuit and should not be generalised without the multi-seed study left to future work. Because the projection discards the records that noise pushes off the physical support, it also acts as an error-mitigation step for this application, at the price of a discard rate that rises from 7.2% to 50.0% over the noise range.

### *G. Limitations and Threats to Validity*

All quantum results are classical simulations at fourteen qubits, and the exact Born vector used for the projection and for the training-curve diagnostics is not available on hardware; the shot-based experiment of Section VIII-F is the hardware-relevant one. The data come from a phenomenological bearing model rather than from a test rig, which is what makes a ground-truth feasibility support and a 24 000-record test set possible, but it also means that the class-conditional indicator distributions are close to Gaussian by construction and favour the GMM; on measured data with multi-modal operating regimes the ordering between the GMM and the kernel-based generators may differ, and the study makes no claim beyond the simulated distribution. The feasibility support was derived from the same physical model that generated the data; a misspecified support would remove real records (the deployment distribution already loses 1.8% of its mass at the sweep extremes) and must be validated against a sample of real records before deployment. The four-indicator space is low-dimensional; the known arguments for quantum generative advantage concern distributions whose structure is hard to capture classically [3], [23], and the present results should be read as showing that a four-indicator bearing space is not such a case. Finally, the membership-inference attack is a black-box frequency attack; stronger white-box attacks on the circuit parameters were not evaluated.

## IX. COMPARATIVE ANALYSIS

Table V places the proposed framework against the published quantum generative studies and the classical augmentation practice for fault diagnosis. Direct numerical comparison with [1], [2], [7], [23] and [25] is not meaningful because those works use different target distributions, so the comparison is made on the dimensions that determine practical use: the kind of data, the sample regime, whether physical admissibility is enforced, whether privacy is measured, whether the classical comparison is controlled, the hardware realism of the evaluation, and the outcome relative to the best classical baseline. The entries for this work are taken from Section VIII.

Three points follow from Table V. First, the present study is the only one that measures privacy leakage of a quantum generator, and it finds the Born machine no more leaky than a Gaussian mixture and much less leaky than a classical energy model of the same size and objective. Second, it is the only one that enforces a physically derived admissibility constraint; the constraint turned out to be generator-agnostic and to be the single most useful ingredient for downstream utility, lifting the QCBM macro-

F1 by 0.15 at $N = 64$ and the cGAN inner-race recall by 0.14 at $N = 128$. Third, its outcome agrees with [23]: on a low-dimensional, approximately Gaussian target a quantum generator does not outperform a well-chosen classical density estimator, and the benefit is confined to the comparison within the family of MMD-trained parametric models. In terms of cost, the QCBM requires 126 parameters and, on hardware, tens of millions of shots per training run, against a few hundred parameters and milliseconds of computation for the GMM; in terms of scalability, the simulator cost grows as $2^n$ while the kernel cost grows as the square of the number of feature cells, so the present evaluation protocol cannot be extended much beyond twenty qubits without hardware. Sustainability is therefore an argument against the quantum generator at this scale, since its training energy is dominated by repeated circuit execution that the classical baselines do not require.

TABLE V COMPARATIVE ANALYSIS WITH EXISTING QUANTUM GENERATIVE AND CLASSICAL AUGMENTATION STUDIES

| Study | Data; constraint / privacy / classical control; hardware realism | Outcome relative to the best classical baseline |
|---|---|---|
| Benedetti *et al.* [1]; Liu and Wang [2] | Bars-and-stripes, GHZ, mixtures, full distribution; none / none / none; 4-qubit trapped-ion device [1], simulation with shot estimator [2] | Not compared |
| Rudolph *et al.* [7] | Handwritten digits, thousands; none / none / shared classical decoder; 8-qubit ion trap | Hybrid; circuit effect not isolated |
| Gili *et al.* [25] | Cardinality-constrained strings, hundreds; combinatorial / none / none; simulation | Generalises; no classical baseline |
| Hibat-Allah *et al.* [23] | Cardinality sets, 20 qubits; combinatorial / none / RNN, transformer, WGAN; simulation | No consistent quantum advantage |
| Zhang *et al.* [21] | Measured vibration spectra, thousands; none / none / classical only | GAN improves imbalanced accuracy |
| This work | Physics-based bearing indicators, 14 qubits, 32–256 records with 5% rare class; physical support Ω / copy ratio and membership AUC / matched C-MMD plus GMM, KDE, cGAN with the same Ω; simulation, 500–2000-shot training, depolarising noise, heavy-hex transpilation | Beats matched C-MMD (TVD 0.52 vs 0.70 at N = 32) and is as private as a GMM; below GMM/KDE on fidelity (masked GMM TVD 0.36 at N = 32) and on utility |

## X. KEY OUTCOMES

*Novel contributions.* The framework is the first to couple a quantum circuit Born machine with a feasibility support derived from the physics of the monitored asset, applied as a soft penalty and as a total-variation-optimal post-selection (Proposition 1), and to embed ordinal condition indicators with a Gray code. The evaluation protocol, which fixes the records, the constraint, the objective and the parameter count across quantum and classical generators and measures privacy leakage, can be reused for other quantum generative claims.

*Engineering findings and significance.* On 32 to 256 bearing records the projected QCBM produces only physically admissible data, improves its own TVD by up to 0.04 and its downstream macro-F1 by up to 0.30 through projection, leaks no more membership information than a Gaussian mixture, lifts rare-class recall to the level of the best classical augmentation, and surpasses a parameter-matched classical energy model at $N \leq 64$; it does not match Gaussian-mixture or kernel-density generators on fidelity, coverage or macro-F1. The feasibility projection is immediately usable with any generator and is the recommended first step for synthetic sensor data in condition monitoring: it removed 8.5% impossible records from a kernel-density generator and 8.2% from a conditional GAN at $N = 64$ while leaving their logistic-regression utility unchanged within 0.01 and improving their random-forest utility. The copy-ratio and membership-AUC audit provides a quantitative release criterion for plants that must share fault data, and the negative comparison guards against premature procurement of quantum resources for a task that classical density estimation already solves.

## XI. CONCLUSION AND FUTURE SCOPE

This paper asked whether a physics-constrained quantum generative model can produce useful synthetic industrial sensor data from a few dozen records. A conditional fourteen-qubit QCBM with a feasibility support derived from a rolling-element bearing model, exact adjoint and finite-shot gradients and depolarising-noise simulation was evaluated on fidelity, diversity, validity, privacy and downstream utility against four classical generators under identical data and constraints. The feasibility projection, proven total-variation optimal, removes every physically impossible record and improves the quantum generator on every criterion, and training from 2000 measurement shots per circuit execution matches the exact gradient. The quantum generator outperforms a parameter-matched classical energy model in the scarcest regime and is as private as a Gaussian mixture, but it is dominated by Gaussian-mixture and kernel-density estimators on the four-indicator bearing space, whose class-conditional structure is close to Gaussian. All findings are simulation-based and rest on a phenomenological signal model; the negative comparison is reported deliberately because it determines where the method should and should not be applied.

*Future scope.* Four directions follow. (i) Validation on measured vibration archives, where multi-modal operating regimes may weaken the Gaussian-mixture advantage and where the feasibility support must be learnt from a digital twin and audited against real records. (ii) Higher-

dimensional discrete targets, such as joint distributions over tens of spectral bands or over sequences of fault states, where classical density estimators lose their inductive advantage and where the known hardness arguments for Born machines apply; the evaluation protocol of this paper transfers unchanged. (iii) Execution on superconducting or trapped-ion hardware with the shot budgets quantified in Section VIII-F, with the projection used as a measurement-level error-mitigation step and with error-aware training under the noise model of Section VIII-F. (iv) Formal privacy guarantees, for example differentially private MMD training of the circuit, and white-box membership attacks on the circuit parameters, to replace the black-box audit used here. A Qiskit transpilation of the circuit for a heavy-hexagon coupling map is included in the reproducibility bundle as a starting point for (iii).

## REFERENCES


[1] M. Benedetti, D. Garcia-Pintos, O. Perdomo, V. Leyton-Ortega, Y. Nam, and A. Perdomo-Ortiz, "A generative modeling approach for benchmarking and training shallow quantum circuits," *npj Quantum Inf.*, vol. 5, Art. no. 45, 2019.

[2] J.-G. Liu and L. Wang, "Differentiable learning of quantum circuit Born machines," *Phys. Rev. A*, vol. 98, Art. no. 062324, 2018.

[3] B. Coyle, D. Mills, V. Danos, and E. Kashefi, "The Born supremacy: Quantum advantage and training of an Ising Born machine," *npj Quantum Inf.*, vol. 6, Art. no. 60, 2020.

[4] C. Zoufal, A. Lucchi, and S. Woerner, "Quantum generative adversarial networks for learning and loading random distributions," *npj Quantum Inf.*, vol. 5, Art. no. 103, 2019.

[5] S. Lloyd and C. Weedbrook, "Quantum generative adversarial learning," *Phys. Rev. Lett.*, vol. 121, Art. no. 040502, 2018.

[6] P.-L. Dallaire-Demers and N. Killoran, "Quantum generative adversarial networks," *Phys. Rev. A*, vol. 98, Art. no. 012324, 2018.

[7] M. S. Rudolph, N. B. Toussaint, A. Katabarwa, S. Johri, B. Peropadre, and A. Perdomo-Ortiz, "Generation of high-resolution handwritten digits with an ion-trap quantum computer," *Phys. Rev. X*, vol. 12, Art. no. 031010, 2022.

[8] M. Cerezo, A. Arrasmith, R. Babbush, S. C. Benjamin, S. Endo, K. Fujii, J. R. McClean, K. Mitarai, X. Yuan, L. Cincio, and P. J. Coles, "Variational quantum algorithms," *Nat. Rev. Phys.*, vol. 3, pp. 625–644, 2021.

[9] J. R. McClean, S. Boixo, V. N. Smelyanskiy, R. Babbush, and H. Neven, "Barren plateaus in quantum neural network training landscapes," *Nat. Commun.*, vol. 9, Art. no. 4812, 2018.

[10] A. Gretton, K. M. Borgwardt, M. J. Rasch, B. Schölkopf, and A. Smola, "A kernel two-sample test," *J. Mach. Learn. Res.*, vol. 13, pp. 723–773, 2012.

[11] T. Jones and J. Gacon, "Efficient calculation of gradients in classical simulations of variational quantum algorithms," 2020, *arXiv:2009.02823*.

[12] V. Bergholm, J. Izaac, M. Schuld, C. Gogolin, S. Ahmed, V. Ajith, M. S. Alam, G. Alonso-Linaje, B. AkashNarayanan, A. Asadi, *et al.*, "PennyLane: Automatic differentiation of hybrid quantum-classical computations," 2018, *arXiv:1811.04968*.

[13] M. Schuld, V. Bergholm, C. Gogolin, J. Izaac, and N. Killoran, "Evaluating analytic gradients on quantum hardware," *Phys. Rev. A*, vol. 99, Art. no. 032331, 2019.

[14] R. B. Randall and J. Antoni, "Rolling element bearing diagnostics—A tutorial," *Mech. Syst. Signal Process.*, vol. 25, no. 2, pp. 485–520, 2011.

[15] W. A. Smith and R. B. Randall, "Rolling element bearing diagnostics using the Case Western Reserve University data: A benchmark study," *Mech. Syst. Signal Process.*, vol. 64–65, pp. 100–131, 2015.

[16] R. Shokri, M. Stronati, C. Song, and V. Shmatikov, "Membership inference attacks against machine learning models," in *Proc. IEEE Symp. Security and Privacy (SP)*, San Jose, CA, USA, 2017, pp. 3–18.

[17] J. Hayes, L. Melis, G. Danezis, and E. De Cristofaro, "LOGAN: Membership inference attacks against generative models," *Proc. Privacy Enhancing Technol.*, vol. 2019, no. 1, pp. 133–152, 2019.

[18] L. Xu, M. Skoularidou, A. Cuesta-Infante, and K. Veeramachaneni, "Modeling tabular data using conditional GAN," in *Advances in Neural Information Processing Systems 32 (NeurIPS)*, Vancouver, Canada, 2019.

[19] A. Alaa, B. van Breugel, E. S. Saveliev, and M. van der Schaar, "How faithful is your synthetic data? Sample-level metrics for evaluating and auditing generative models," in *Proc. 39th Int. Conf. Machine Learning (ICML)*, Baltimore, MD, USA, 2022, pp. 290–306.

[20] N. V. Chawla, K. W. Bowyer, L. O. Hall, and W. P. Kegelmeyer, "SMOTE: Synthetic minority over-sampling technique," *J. Artif. Intell. Res.*, vol. 16, pp. 321–357, 2002.

[21] W. Zhang, X. Li, X.-D. Jia, and Q. Ding, "Machinery fault diagnosis with imbalanced data using deep generative adversarial networks," *Measurement*, vol. 152, Art. no. 107377, 2020.

[22] J. Preskill, "Quantum computing in the NISQ era and beyond," *Quantum*, vol. 2, p. 79, 2018.

[23] M. Hibat-Allah, M. Mauri, J. Carrasquilla, and A. Perdomo-Ortiz, "A framework for demonstrating practical quantum advantage: Comparing quantum against classical generative models," *Commun. Phys.*, vol. 7, Art. no. 68, 2024.

[24] A. Javadi-Abhari, M. Treinish, K. Krsulich, C. J. Wood, J. Lishman, J. Gacon, S. Martiel, P. D. Nation, L. S. Bishop, A. W. Cross, B. R. Johnson, and J. M. Gambetta, "Quantum computing with Qiskit," 2024, *arXiv:2405.08810*.

[25] K. Gili, M. Hibat-Allah, M. Mauri, C. Ballance, and A. Perdomo-Ortiz, "Do quantum circuit Born machines generalize?," *Quantum Sci. Technol.*, vol. 8, Art. no. 035021, 2023.

[26] D. P. Kingma and M. Welling, "Auto-encoding variational Bayes," in *Proc. 2nd Int. Conf. Learning Representations (ICLR)*, Banff, Canada, 2014.

[27] I. Goodfellow, J. Pouget-Abadie, M. Mirza, B. Xu, D. Warde-Farley, S. Ozair, A. Courville, and Y. Bengio, "Generative adversarial nets," in *Advances in Neural Information Processing Systems 27 (NIPS)*, Montreal, Canada, 2014, pp. 2672–2680.

[28] M. S. Rudolph, S. Lerch, S. Thanasilp, O. Kiss, O. Shaya, S. Vallecorsa, M. Grossi, and Z. Holmes, "Trainability barriers and opportunities in quantum generative modeling," *npj Quantum Inf.*, vol. 10, Art. no. 116, 2024.

[29] J. Antoni, "Fast computation of the kurtogram for the detection of transient faults," *Mech. Syst. Signal Process.*, vol. 21, no. 1, pp. 108–124, 2007.

[30] D. P. Kingma and J. Ba, "Adam: A method for stochastic optimization," in *Proc. 3rd Int. Conf. Learning Representations (ICLR)*, San Diego, CA, USA, 2015.