

Supplementary Material

This appendix supplements the main paper with expanded content, starting with a detailed background on gate-based quantum computing, covering its foundations in quantum physics and applications in machine learning and INRs (Sec. A). It explores practical implementations supported by QVF such as image inpainting (Sec. B) and shape completion (Sec. C), complemented by additional visualizations of 3D geometries (Sec. D) and the quantum ansatz architecture used in primary experiments (Sec. J). The appendix further investigates parameter scaling dynamics (Sec. E), outlines the full algorithmic protocol (Sec. F), and expands on related literature (Sec. G). Theoretical contributions include a proof demonstrating the efficacy of the ansatz in QVF, which feature bounded haar randomness, for enhanced learning (Sec. H), alongside an analysis linking quantum circuit measurements to Bayesian inference frameworks (Sec. I).

A Background

A.1 Preliminaries of Gate-based Quantum Computing

Qubits. The fundamental information blocks of a quantum processing unit (QPU) are qubits, i.e., the analogues of bits in classical computing. Unlike classical bits deterministically representing one possible state (0 or 1), qubits can statistically represent two distinct information states at the same time, denoted in the bra-ket notation as $|0\rangle$ and $|1\rangle$.

Superposition is a fundamental property distinguishing qubits from bits: It grants qubits the capacity to exist in a combinatorial state $|\psi\rangle$ of $|0\rangle$ and $|1\rangle$ such that:

$$|\psi\rangle = \alpha |0\rangle + \beta |1\rangle, \quad (13)$$

with $\alpha, \beta \in \mathbb{C}$ and $|\alpha|^2 + |\beta|^2 = 1$. Qubit states $|\psi\rangle$ can be visualised on Bloch spheres (see Fig. 10) or expressed in a vector form:

$$|0\rangle = \begin{bmatrix} 1 \\ 0 \end{bmatrix}, |1\rangle = \begin{bmatrix} 0 \\ 1 \end{bmatrix}, |\psi\rangle = \alpha |0\rangle + \beta |1\rangle = \begin{bmatrix} \alpha \\ \beta \end{bmatrix}. \quad (14)$$

Measurement in quantum mechanics inherently adopts a statistical approach to extract numerical information. For a qubit state $|\psi\rangle = \alpha |0\rangle + \beta |1\rangle$ measured with operator \hat{O} (that must be Hermitian, i.e., $\hat{O}^\dagger = \hat{O}$), this implies probabilities $|\alpha|^2$ and $|\beta|^2$, respectively, for measuring the information (i.e., eigenvalue of the measurement operator \hat{O}) stored in states $|0\rangle$ and $|1\rangle$:

$$\hat{O} |0\rangle = \kappa |0\rangle \text{ and } \hat{O} |1\rangle = \delta |1\rangle, \quad (15)$$

where κ and δ are eigenvalues of the measurement operator $|O\rangle$. The key aspect of measurement is the phenomenon known as wave function collapse, i.e., the projective measurement causes $|\psi\rangle$ to collapse to the operator's eigenstate, $|0\rangle$ or $|1\rangle$, conditioned on the measurement, i.e., κ or δ .

Entanglement further distinguishes quantum from classical computing. In the classical case, information stored in bits is independent, i.e., measuring one bit does not affect others. In the quantum realm, qubits can be highly correlated, exhibiting entanglement such that the information of one qubit can be interrelated with another, despite possible physical distance between them. For instance, a general information state of a 2-qubit system $|\psi\rangle_2$ can be expressed as:

$$|\psi\rangle_2 = a |00\rangle + b |01\rangle + c |10\rangle + d |11\rangle, \quad (16)$$

with $a, b, c, d \in \mathbb{C}$ such that $|a|^2 + |b|^2 + |c|^2 + |d|^2 = 1$. The 2-qubit system is considered entangled if $|\psi\rangle_2$ cannot be expressed as a tensor product of two qubits $|\psi\rangle_{a1}$ and $|\psi\rangle_{a2}$, indicating that their information cannot be independently measured without disturbing each other, i.e.,

$$|\psi\rangle_2 \neq |\psi\rangle_{a1} \otimes |\psi\rangle_{a2}. \quad (17)$$

Rotation Operators. The operators responsible for rotating quantum states $|\psi\rangle$ of qubits along x, y, z axes on a Bloch sphere are referred to as rotation operators. Any single qubit operator \hat{R} can be expressed as a combination of such rotation operators $\hat{R}_x, \hat{R}_y, \hat{R}_z$, i.e., $\hat{R}(\theta, \tau, \gamma) = \hat{R}_x(\theta) \hat{R}_y(\tau) \hat{R}_z(\gamma)$ with angles θ, τ and γ :

$$\hat{R}_x(\theta) = \begin{bmatrix} \cos(\frac{\theta}{2}) & -i \sin(\frac{\theta}{2}) \\ -i \sin(\frac{\theta}{2}) & \cos(\frac{\theta}{2}) \end{bmatrix}, \quad \hat{R}_y(\tau) = \begin{bmatrix} \cos(\frac{\tau}{2}) & -\sin(\frac{\tau}{2}) \\ \sin(\frac{\tau}{2}) & \cos(\frac{\tau}{2}) \end{bmatrix}, \quad \hat{R}_z(\gamma) = \begin{bmatrix} e^{-i\frac{\gamma}{2}} & 0 \\ 0 & e^{i\frac{\gamma}{2}} \end{bmatrix}. \quad (18)$$




Figure 10: Bloch sphere visualisation of qubit states. Qubit 0: $|\psi\rangle = \frac{1}{\sqrt{2}}(|0\rangle + |1\rangle)$, qubit 1: $|\psi\rangle = |1\rangle$.

587 The Pauli operators $\hat{X}, \hat{Y}, \hat{Z}$ represent specific instances of above rotation operators, i.e. rotations by π radians
 588 along the x, y, z -axes, respectively. These operators can also be expressed as matrices in the computational basis
 589 $|0\rangle, |1\rangle$ as follows:

$$\hat{X} = \begin{bmatrix} 0 & 1 \\ 1 & 0 \end{bmatrix}, \hat{Y} = \begin{bmatrix} 0 & -i \\ i & 0 \end{bmatrix}, \hat{Z} = \begin{bmatrix} 1 & 0 \\ 0 & -1 \end{bmatrix}. \quad (19)$$

590 **The Schrödinger Equation.** Quantum computing involves the manipulation of information according to the
 591 principles of quantum mechanics, with its foundation rooted in the time-dependent Schrödinger equation:

$$i\hbar \frac{d}{dt} |\psi(t)\rangle = \hat{H}(t) |\psi(t)\rangle, \quad (20)$$

592 where \hbar is Planck's constant, and $|\psi(t)\rangle$ and $|\psi(0)\rangle$ are the quantum states after and before evolution, respec-
 593 tively; \hat{H} is the Hamiltonian operator of the quantum system. Therefore, the evolution of quantum states can be
 594 described by the following relationship:

$$|\psi(t)\rangle = \hat{T} e^{-\frac{i}{\hbar} \int_0^t \hat{H}(t) dt} |\psi(0)\rangle, \quad (21)$$

595 with \hat{T} denoting the time ordering operator. This simplifies to $e^{-\frac{it}{\hbar} \hat{H}} |\psi(0)\rangle$ for time-independent \hat{H} . Using a
 596 more compact notation, the Schrödinger equation can also be equivalently written as:

$$|\psi(t)\rangle = \hat{U}(\hat{H}, t) |\psi(0)\rangle, \text{ with} \quad (22)$$

597

$$\hat{U}(\hat{H}, t) = e^{-\frac{it}{\hbar} \hat{H}}. \quad (23)$$

598 To perform rotation operations on qubits, the system Hamiltonian \hat{H} can be set to $E\hat{\sigma}$, with $\hat{\sigma} \in \{\hat{X}, \hat{Y}, \hat{Z}\}$. By
 599 substituting $\eta = 2Et/\hbar$, we arrive at:

$$\hat{U}(\hat{H}, t) = e^{-\frac{it}{\hbar} \hat{H}} = e^{-\frac{i\eta}{2} \hat{\sigma}} = \hat{R}_{\sigma}(\eta). \quad (24)$$

600 A.2 Preliminaries of INR

601 For completeness, we provide a brief review of INR. INR models continuous signals—such as images, 3D
 602 geometries, or physical fields—as neural networks. Unlike traditional explicit representations that store data
 603 discretely (e.g., pixels or voxels), INR defines a signal as a continuous function $f_{\theta} : \mathbb{R}^d \rightarrow \mathbb{R}^k$, where a neural
 604 network parameterized by θ maps spatial or temporal coordinates to corresponding signal values, enabling
 605 resolution-independent, memory-efficient, and differentiable representations, making INR particularly valuable
 606 for tasks requiring interpolation, super-resolution, or solving inverse problems. The network is trained to
 607 minimize the reconstruction loss over sampled coordinates $\mathcal{L}(\theta) = \sum_{\mathbf{x} \in \mathcal{X}} \|f_{\theta}(\mathbf{x}) - \mathcal{S}(\mathbf{x})\|^2$ where \mathcal{X} denotes
 608 the domain of interest. INR architectures typically address spectral bias—the tendency to favor low-frequency
 609 components—through positional encoding or periodic activation functions such as SIREN (45).

610 A.3 Quantum Machine Learning

611 The potential of quantum computing to enhance machine learning algorithms leads to the emergence of quantum
 612 machine learning (QML) (40), a discipline employing quantum mechanical phenomena to tackle classically
 613 intractable learning problems through enhanced computational paradigms (40). Central to QML are: 1) a
 614 feature map, which encodes classical input data into quantum states and 2) a variational ansatz, which performs
 615 quantum transformation on the quantum states. PQC have been shown to be asymptotic universal function
 616 approximators (5; 41). Several standardized QML algorithms have been explored, including quantum principal
 617 component analysis (28), quantum support vector machines (37), quantum Boltzmann machines (2), and quantum
 618 k-means clustering (24).

619 **Feature Map.** Integrating classical Euclidean data \mathbf{x} into quantum computational frameworks necessitates a
 620 non-trivial mapping to quantum states $|\psi(\mathbf{x})\rangle$ in a Hilbert space \mathcal{H} . Several established encoding techniques
 621 exist, including *basis encoding*, *amplitude encoding*, *Hamiltonian evolution encoding*, with each presenting
 622 distinct trade-offs in qubit efficiency and circuit depth complexity. However, the determination of optimal
 623 encoding schemes remains an open research challenge, as the relationship between encoding fidelity $\mathcal{F}(\mathbf{x}) =$
 624 $|\langle \psi_{\text{ideal}}(\mathbf{x}) | \psi_{\text{encoded}}(\mathbf{x}) \rangle|^2$, resource requirements, and task-specific performance metrics (e.g., classification
 625 accuracy or function approximation error ϵ) remains poorly characterized across different problem domains.

626 **Variational Ansatz.** Quantum evolution of classical information embedded in states $|\psi(x)\rangle$ requires parameter-
 627 ized unitary ansatz $\hat{U}(\theta) \in \mathbb{C}^{2^n \times 2^n}$ acting on n -qubit systems. Physically, the ansatz is constructed through
 628 sequential composition of such unitary transformations, formally expressed as $\hat{U}(\theta) = \mathcal{T} \left(\prod_{i=1}^t \hat{U}_i(\theta_i) \right)$,
 629 where \mathcal{T} denotes the time-ordering operator governing gate sequence implementation. This induces a Hilbert

space transformation $\hat{U}(\theta) : \mathcal{H} \rightarrow \mathcal{H}$ that maps input states to processed output states through the operation $|\phi(x, \theta)\rangle = \hat{U}(\theta) |\psi(x)\rangle$.

Measurement. Quantum computation culminates in statistical data extraction from evolved quantum states $|\phi(x)\rangle$ through projective measurements using Hermitian observables \hat{O} , where the computational output is formally defined as the expectation value: $V(x) = \langle \phi(x) | \hat{O} | \phi(x) \rangle$. Measurements collapse quantum states according to the Born rule, thereby restricting access to the embedded classical information to statistical estimators derived from repeated measurements. The choice of observable \hat{O} fundamentally governs both the information-theoretic capacity of the measurement protocol and its computational complexity.

Training a Variational Ansatz. Instead of constructing a computational graph and performing backpropagation, training quantum circuits involves only forward evaluations (50). To minimize a measurement-dependent cost function $\mathcal{L}(\theta)$, the exact gradients $\nabla \mathcal{L}(\theta)$ can be evaluated through quantum circuit evaluations at shifted parameters $\theta \pm \frac{\pi}{2} \mathbf{e}_i$ for basis vectors \mathbf{e}_i , expressed as:

$$\partial_{\theta_i} \mathcal{L} = \frac{1}{2} \left[\mathcal{L} \left(\theta_i + \frac{\pi}{2} \right) - \mathcal{L} \left(\theta_i - \frac{\pi}{2} \right) \right]. \quad (25)$$

This technique, i.e. parameter-shift rule, exploits the trigonometric structure of unitary gate generators \hat{G}_i (where $\hat{U}_i(\theta_i) = e^{-i\theta_i \hat{G}_i}$) to enable hardware-compatible gradient estimation without numerical approximation or persistent circuit memory - a critical advantage over classical backpropagation that requires differentiable computational graphs.

A.4 Barren Plateaus

Training a variational ansatz $\hat{S}(\theta)$ is fundamentally constrained by the *barren plateau* phenomenon, where random parameter initialization induces exponential vanishing of cost function gradients across Hilbert space. As formally demonstrated by McClean et al. (30) through concentration of measure analysis:

“...for a wide class of reasonable parametrized quantum circuits, the probability that the gradient along any reasonable direction is non-zero to some fixed precision is exponentially small as a function of the number of qubits.”

This observation is also known as *barren plateau*, which can be expressed mathematically for a system with n qubits as:

$$\mathbb{E}_w[\partial_w L(w)] = 0, \quad \text{Var}_w[\partial_w L(w)] \in O\left(\frac{1}{\nu^n}\right), \quad \nu > 1, \quad (26)$$

where ν characterizes the circuit’s entangling capacity. The variance bound’s scaling establishes that gradient estimators require $\mathcal{O}(\nu^n)$ measurement samples to maintain constant precision - an exponential resource overhead that renders practical optimization infeasible for $n \gg 1$. This poses challenges, particularly for gradient-based learning. Identified factors contributing to this phenomenon include observable locality (10; 46), specific noise models (48), and ansatz close to a 2-design, i.e., matching Haar random unitaries up to the second moment (30; 21). Those highlight the importance of selecting appropriate initialization protocols, quantum ansatz designs and observables.

B Application: Image Inpainting

QIREN (53) and 3D-QAE (36) are constrained by their reliance on fixed latent representations or rigid interpolation mechanisms, thereby not capable to reconstruct complete, coherent outputs from partial or corrupted inputs. QVF addresses such limitations by conditioning the quantum circuit topology on both the query point and a dynamic latent space, benefiting applications such as image inpainting. Given images with occluded or corrupted pixels, the circuit identifies in the latent space a vector that minimizes the discrepancy between the predicted multi-dimensional properties learned by the quantum circuit and the observed noisy values. The optimal latent vector conditioned quantum circuit enables recovery of missing field properties across the whole domain. Empirically, we masked out half of the image pixels and reconstructed the complete images via the protocol. Representative results of image inpainting with QVF are visualized in Fig. 11, demonstrating that QVF can deliver promising performance and accurately recover images even under such extreme sparsity, positioning QVF as a promising quantum circuit architecture for such tasks.




Figure 11: Image inpainting results with QVF.




Figure 12: Shape completion from partial and noisy input depth maps using QVF; α is the noise ratio.

C Application: Shape Completion from Partial and Noisy Depth Maps

With the principle similar to image inpainting, QVF can perform tasks such as completing 3D geometries while existing quantum models can not. We adopt similar empirical settings as for images by cropping half of the samples along the depth (see Fig. 8-(b)) and consider the effects of zero-mean Gaussian noise across different perturbation ratios α to the depth map. Shape completion performance is quantified across incremental perturbation ratios, parameterized as $\alpha \in \{0, 0.005, 0.01, 0.02, 0.03\}$, where $\alpha = 0$ corresponds to an idealized noise-free completion scenario. Quantitative analysis reveals a monotonic decline in reconstruction fidelity with increasing noise perturbations, as evidenced by progressive geometric distortions and surface irregularities in Fig. 12.

D Additional 3D Geometry Visualizations

In conjunction with the quantitative results summarized in Tab. 3, Fig. 13-(a) provides a qualitative comparison of the reconstructed 3D geometries, contrasting the baseline—the classical architectural component in our model—with QVF. The baseline’s numerically elevated loss values correlate with visual structural discontinuities, exemplified by the fragmented sofa leg, underscoring its propensity for topological inconsistencies during reconstruction. QVF, instead, demonstrates enhanced structural coherence, generating topologically intact geometries devoid of visible artifacts, as evidenced by its preservation of fine-grained features. Further geometric analysis, illustrated via color-encoded per-surface Hausdorff distance distributions in Fig. 13-(b), reveals systematic geometric deviations for the baseline (top) and ours (bottom), corroborating its geometric fidelity. An interesting observation is that the color distributions between models align well, meaning that QVF inherits the representative expressiveness of the classical component but enhances it due to inherent spectral connections to the quantum circuit.




Figure 13: (a) Comparison of geometric representations using QVF and the classical model; ground truth is presented at the bottom; (b) representation fidelity visualized via color-encoded Hausdorff distance map: colors represent the distance to the ground truth. The rendered image employs a color gradient (blue>green>yellow>red) to indicate descending Hausdorff distance levels.

E Parameter Scaling Analysis

Complementing the visualization of model parameters in Fig. 9, we present a rigorous mathematical analysis of parameter scaling in QVF. The classical component, which we model as a dense neural network for inferring quantum system energy spectra, exhibits quadratic parameter growth $\mathcal{O}(p^2)$ with respect to hidden dimension p - a direct consequence of its fully connected architecture requiring pairwise neuron interactions. In contrast, the quantum evolution parameters demonstrate parametrically efficient linear scaling $\mathcal{O}(n\mathcal{J})$ for an n -qubit circuit of depth \mathcal{J} . This fundamental interdependence necessitates balanced co-design: while quantum circuit

707 scaling enhances expressivity through Hilbert space exploration, the operational fidelity critically depends on
 708 the classical network’s capacity to provide accurate spectra. Optimal system performance requires concurrent
 709 optimization of classical representation capacity and quantum resource allocation.

710 F Algorithmic Protocol

We provide the complete training protocol of QVF in Alg. 1.

Algorithm 1 QVF Training Protocol

Input: Training dataset $\mathcal{D} = \{(x_i, s_i)\}_{i=1}^N$; number of qubits n ; epochs N_{epoch} ; measurement shots N_{shot} ; parameters θ_q, θ_c ; inverse temperature β .

for epoch = 1 **to** N_{epoch} **do**

Classical Inference:

- Compute energy spectrum $\mathbf{E}(\mathbf{x}; \theta_c)$ via classical model.
- Evaluate Gibbs distribution:

$$P_i = \frac{e^{-\beta E_i}}{Z}, \text{ where } Z = \sum_{j=1}^{2^n} e^{-\beta E_j}.$$

Quantum State Preparation:

 Initialize $\hat{\rho}_0 = \sum_{i=1}^{2^n} P_i |i\rangle \langle i|$.

Quantum Evolution:

 Apply ansatz $\hat{S}(\theta_q) = \prod_{\ell=1}^J e^{-i\theta_{q,\ell} \hat{H}_\ell}$ to obtain:

$$\hat{\rho}(\theta_q) = \hat{S}(\theta_q) \hat{\rho}_0 \hat{S}^\dagger(\theta_q)$$

Measurement & Observables:

 Estimate $\langle \hat{O}_k \rangle = \text{Tr}[\hat{\rho}(\theta_q) \hat{O}_k]$ for $k = 1, \dots, K$.

Gradient Computation:

- Quantum: $\partial \langle \hat{O}_k \rangle / \partial \theta_{q,\ell}$ via parameter-shift rule.
- Classical: $\nabla_{\theta_c} \mathcal{L}$ via chain rule through $\mathbf{E}(\mathbf{x}; \theta_c)$.

Parameter Update:

 Adam optimizer step with learning rate η :

$$\theta \leftarrow \theta - \eta (\nabla_{\theta_c} \mathcal{L}, \nabla_{\theta_q} \mathcal{L}).$$

end for

Output: Optimized parameters θ_q^*, θ_c^* .

711

712 G Extended Literature Discussion

713 Recent frameworks, including QIREN (53) and related works (42; 1; 27; 20; 23), adopt a unified architectural
 714 paradigm—termed the sandwich structure—in which a quantum circuit is embedded between classical prepro-
 715 cessing and postprocessing layers. Such a framework positions PQCs as trainable Fourier feature encoders,
 716 theoretically capable of projecting classical inputs into high-dimensional spectral representations under optimal
 717 conditions, subsequently processed by classical neural networks. While this design draws conceptual parallels
 718 to classical Fourier Feature Networks (FFNs) by substituting hidden layers with quantum circuits to explore
 719 quantum-enhanced feature spaces, it introduces fundamental conceptual and operational trade-offs that need to
 720 be carefully considered:

721 **Over-reliance on Post-processing.** The quantum circuit is constrained solely to be a feature generator, ne-
 722 cessitating extensive post-processing by classical neural networks to extract meaningful information from
 723 the generated features. This dependency challenges claims of inherent quantum advantage, as the quantum
 724 subsystem does not directly contribute to the final inference step.

725 **Blurring the Divide: Quantum Probabilism vs. Classical Determinism.** Inference from a quantum circuit is
 726 naturally probabilistic by aggregating statistical projective measurements. Sandwiched structures, however, do
 727 not consider such a probabilistic nature but necessitate single-shot estimations for the inference of the quantum
 728 circuit before classical post-processing. This collapses the quantum subsystem’s probabilistic nature into a




Figure 15: Visualization of ansatz designs and their exemplary induced traversable quantum states within Hilbert space: strongly entangled ansatz (left) and QVF (right). Traversable states of different circuit ansatz are visualized on the bottom right for both ansatzes.

deterministic feature map, as the final output becomes a parametric function of estimated expectation values rather than a direct manifestation of quantum measurement statistics.

H Variational Ansatz with Bounded Haar Randomness Improves Learning

We explain in more detail how our variational ansatz design featuring bounded Haar randomness could better preserve the gradient, which is crucial for inferring precise energy spectra. As a review, for quantum circuits with a generic parameterized ansatz, the gradient expectation value asymptotically scales to zero as the qubit count increases, a consequence of concentration in high-dimensional Hilbert spaces. The vanishing gradient phenomenon, i.e. barren plateaus, is thus governed by the variance decay rate, which determines optimization tractability. We quantify this phenomenon by analyzing the gradient variance Var_{grad} :

$$\text{Var}_{\text{grad}} = E_{\theta \sim U([0, 2\pi])} \left[\text{Var} \left(\left\{ \frac{\partial}{\partial \theta_k} \langle \hat{H} \rangle_{\theta} \right\}_{k=1}^M \right) \right]. \quad (27)$$

We empirically evaluate our variational ansatz across an increasing number of qubits and benchmark against the established strongly entangled ansatz initially proposed in (39), and later adopted in QIREN. The results are shown in Fig. 14, confirming our ansatz design better preserves the gradients — a critical factor for our QVF design. Such behavior, as demonstrated by McClean et al. (30), can also be reflected via a higher KL divergence between the fidelity distribution of the variational ansatz and the Haar-random ensemble; see Tab. 4. Specifically, when the traversable quantum states of the variational ansatz closely resemble Haar-random states, its parameter landscape becomes flat, leading to vanishing gradients and poor optimization. Concurrently, bounding the Haar randomness of the ansatz would not impact the expressivity of the circuit, as Z-basis measurement probabilities remain invariant to the real-subspace constraint, as they solely depend on magnitudes of state amplitudes. We evaluate how the trainability disparity contributes to performance on learning image representations. Quantification results are recorded in Tab. 4.

I Connection of PQCs to Bayesian Inference

As quantum circuits are inherently probabilistic models, they share conceptual parallels with Bayesian inference frameworks. In Bayesian neural networks (BNNs), probabilistic outputs emerge from parameters governed by prior distributions, with training focused on maximizing the conditional likelihood of observed data labels while implicitly updating a posterior distribution over the parameters. For PQC, while they leverage deterministic parameters, they exhibit probabilistic outputs due to the stochastic nature of quantum measurements, which—under sufficiently large shot counts—approximate Gaussian distributions in accordance with the CLT. While the probabilistic outputs of PQCs permit an interpretative lens rooted in Bayesian principles, their training does not inherently involve posterior inference over parameters unless explicitly cast within a Bayesian formalism. This distinction underscores that the Bayesian interpretation of PQCs arises from their measurement statistics rather than an intrinsic probabilistic parameter space.

J Detailed Variational Ansatz Visualization

We provide a visualization of: 1) the QVF ansatz and 2) the widely-adapted strongly-entangled ansatz (39) in Fig. 15. We highlight the circuit structures via their parametrized single-qubit rotations and inter-qubit

	Baseline (39)	Ours
KL divergence	0.0491	0.2392
MSE ($1e^{-3}$) ↓	1.66	1.02

Table 4: Performance comparison.




Figure 14: Gradient variance across circuit of different scales.

768 entanglement patterns. Traversable quantum states are visualized on the Bloch sphere by sampling the ansatz. As
769 expected, our circuit design confines generated quantum states to a circular subset of the Bloch sphere, whereas
770 quantum states evolved by the entangling ansatz are broadly distributed across the complete Bloch sphere surface.
771 Such visualizations complement our observation in Tab. 4, where the ansatz in QVF achieves an elevated KL
772 divergence with Haar random states, leading to stronger gradient magnitudes and improved trainability.