

Algebraic and Topological Limits of Unitary-Input Data Re-Uploading

Mahtab Yaghubi Rad and Vedran Dunjko
Applied Quantum Algorithms (aQa), Leiden University

Abstract

Data re-uploading circuits are a common way to turn an input into a function realized by a parametrized quantum circuit. For scalar data encoded through commuting generators, the resulting model is usually analyzed as a truncated Fourier series: expressivity is controlled by accessible frequencies and reachable Fourier coefficients. We study the non-Abelian counterpart in which the uploaded datum is an unknown unitary matrix $U \in \text{U}(N)$, queried repeatedly between trainable unitary layers and read out through a Born probability. This setting is natural for quantum data, coherent access to unknown transformations, and process-level learning, but its function space is not the naive space of polynomial functions in matrix entries. Physical functions must be invariant under global phase, and polynomial expressions must be interpreted in the coordinate ring obtained after imposing the unitary relations. We formulate this reference space, introduce a balanced degree filtration, and prove three main results. First, depth is algebraically rigid: every positive-depth output has genuine balanced degree exactly equal to the number of uploads. Second, at fixed depth the reachable coefficient image is semi-algebraic and nowhere dense in the ambient finite-dimensional coefficient space. This is stronger than a measure-zero statement: even the closure of the depth- L image contains no open subset of coefficient targets. Third, even unbounded depth is not uniformly universal: a topological range-pinning obstruction separates the full raw model from valid $[0, 1]$ -valued center-invariant targets. As a showcase for the image-geometry tools developed here, we show how semi-algebraicity turns dimension-counting intuition for Boson Sampling Born Machines into a rigorous non-universality statement. The semi-algebraic image geometry developed in this paper provides a novel tool for PQC expressivity theory: it analyzes the intrinsic, stratified geometry of reachable model images, rather than only the local dimension of a particular parametrization.

Problem and model

Parametrized quantum circuits used in quantum machine learning are often judged by whether increasing depth and data re-uploading enlarges the class of functions they can learn. For commuting encodings $E(x) = \prod_j e^{-ix_j H_j}$, the effective encoded phases live on a torus, the model output is a trigonometric polynomial, and the architecture specifies a Fourier spectrum together with constrained coefficient sets [1, 2, 3]. This picture underlies many expressivity and universality intuitions for quantum feature maps and variational models [6].

The present work studies the structurally different case in which the input itself is a general unitary transformation. This is the relevant mathematical abstraction when the data are quantum processes or structured unitary operators supplied coherently to a circuit. It also exposes a limitation of importing Fourier intuitions from scalar re-uploading: for $U \in \text{U}(N)$ the input space is a compact non-Abelian group, and the appropriate harmonic language is Peter–Weyl theory rather than ordinary Fourier analysis. We now fix the raw unitary-input re-uploading architecture studied here.

Definition 1. Let $\mathcal{H}_L$ denote the depth- L squared-amplitude model family

$$\mathcal{H}_L = \{ f_{L,W}(U) = |\langle 0|W_L U \cdots W_1 U W_0|0\rangle|^2, \\ f_{L,W} : \text{U}(N) \rightarrow \mathbb{R}, \quad W \in \text{U}(N)^{L+1} \}. \quad (1)$$

Here the same unknown U acts on the full N -dimensional Hilbert space and is queried L times; no diagonal, local, or commuting structure is assumed. The output is the Born return probability of measuring $|0\rangle$ after the coherent query circuit. The expressivity question is therefore: which real-valued functionals of an unknown transformation can be represented before any question of training, data access, or sample complexity is posed?

The first obstruction is that the target space has to be specified correctly. Since U and zU represent the same physical transformation up to global phase and Eq. (1) is invariant under $U \mapsto zU$, the model should be compared with center-invariant functions, equivalently functions on $\text{PU}(N)$. Moreover, different formal polynomials may agree after restriction to $\text{U}(N)$ because $UU^\dagger = U^\dagger U = I$. We therefore work in the coordinate ring of polynomial functions on the unitary domain and use the balanced filtration

$$\mathcal{F}_0 \subseteq \mathcal{F}_1 \subseteq \cdots \subset R(\text{U}(N), \mathbb{C})^Z,$$

where $\mathcal{F}_L$ is the image of formal polynomials of balanced bi-degree (L, L) after quotienting by the unitary relations. This is the non-Abelian analogue of choosing a Fourier-degree cutoff, but it is sensitive to the algebraic identities of the unitary group.

Algebraic rigidity of depth

The amplitude in Eq. (1) is homogeneous of degree L in the entries of U , while the Born probability is formally bi-homogeneous of degree (L, L) in (U, U^*) . This formal observation is not enough: on $\text{U}(N)$, expressions such as $\sum_k U_{ik} \overline{U_{jk}} = \delta_{ij}$ can lower the formal, apparent degree. Here “genuine” means degree after quotienting by the unitary relations: f has genuine balanced degree L if $f \in \mathcal{F}_L$ but $f \notin \mathcal{F}_{L-1}$, equivalently $[f]_L \neq 0$ in $\mathcal{F}_L/\mathcal{F}_{L-1}$.

Theorem 2. *Let $\bigcup_{l \geq 0} \mathcal{F}_l$ be the balanced polynomial filtration. Then for every $N \geq 2$, $L \geq 1$, $W_0, \dots, W_L \in \text{U}(N)$, the function*

$$U \mapsto |\langle 0 | W_L U W_{L-1} U \cdots W_1 U W_0 | 0 \rangle|^2 \quad (2)$$

has genuine balanced degree exactly L with respect to this filtration. Equivalently,

$$[f_{L,W}]_L \neq 0 \quad \text{as an element of } \mathcal{G}_L := \mathcal{F}_L/\mathcal{F}_{L-1}. \quad (3)$$

The last display says that the class of $f_{L,W}$ remains nonzero after all functions of filtration level at most $L-1$ have been identified with zero.

Lemma 1 (Degree rigidity). *For the fixed-depth model family $\mathcal{H}_L = \{f_{L,W} : \text{U}(N) \rightarrow \mathbb{R} \mid W \in \text{U}(N)^{L+1}\}$, the following properties hold for all L :*

1. $\mathcal{H}_L \subseteq \mathcal{F}_L$;
2. $\mathcal{H}_L \cap \mathcal{F}_{L-1} = \emptyset$.

Thus, depth is not merely adding parameters to a previous model class. Since $\mathcal{H}_k \subseteq \mathcal{F}_k \subseteq \mathcal{F}_{L-1}$ for $k < L$, while $\mathcal{H}_L \cap \mathcal{F}_{L-1} = \emptyset$, functions realized at exact depth L never coincide with lower-depth functions. In particular the usual nested expressivity intuition fails: increasing the number of uploads moves the model into a different genuine-degree sector rather than refining all earlier sectors.

Coefficient geometry at fixed depth

Let $\mathcal{B}_{\leq L} = \{\varphi_1, \dots, \varphi_{D_{N,L}}\}$ be an adapted basis of $\mathcal{F}_L$, meaning that its first $\dim \mathcal{F}_{L-1}$ elements form a basis of $\mathcal{F}_{L-1}$, and write $f_{L,W} = \sum_j c_j(W) \varphi_j$ with unique coefficients $c_j(W) \in \mathbb{C}$. Viewing complex coefficients through real and imaginary parts, the reachable set at depth L is

$$\mathcal{C}_{L,\mathcal{B}}(\mathcal{H}_L) = \{(\text{Re } c_j(W), \text{Im } c_j(W))_{j=1}^{D_{N,L}} : W \in \text{U}(N)^{L+1}\} \subseteq \mathbb{R}^{2D_{N,L}}.$$

For PQC's this image need not be a smooth manifold: redundant parametrizations, phase symmetries, self-intersections, and singularities are typical. Real algebraic geometry gives the right language for such reachable sets. We prove that the coefficient image is semi-algebraic, hence has an intrinsic dimension, a finite smooth stratification, and closure with the same dimension.

Theorem 3 (Nowhere denseness of the coefficient image). *For all $N \geq 2$ and $L \geq 1$, the coefficient image of the depth- L squared-amplitude model on $U(N)$,*

$$\mathcal{C}_{L,B}(\mathcal{H}_L) \subseteq \mathbb{R}^{2D_{N,L}}, \quad (4)$$

is a semi-algebraic set. It has empty interior and is nowhere dense in the Euclidean topology.

This turns the intuitive idea of parameter-counting into a rigorous obstruction for nonlinear quantum model maps: at fixed depth, the coefficient image contains no open subset of $\mathcal{F}_L$, hence no open family of coefficient targets in $\mathcal{F}_L$ can be approximated arbitrarily well.

This semi-algebraic obstruction applies beyond coefficient spaces. More generally, if $\Phi : \mathcal{P} \rightarrow \mathcal{T}$ is a semi-algebraic map from a semi-algebraic parameter space into a semi-algebraic target and $\dim_{\text{sa}} \mathcal{P} < \dim_{\text{sa}} \mathcal{T}$, then $\Phi(\mathcal{P})$ is nowhere dense in $\mathcal{T}$. For a collision-free k -photon, m -mode Boson Sampling Born Machine [4, 5], outcomes are k -subsets $S \subseteq [m]$ and

$$q_{m,k,U}(S) = |\text{Per}(U_{S,[k]})|^2.$$

Hence the model is a polynomial, or after normalization rational, semi-algebraic map $\Phi_{m,k} : U(m) \rightarrow \Delta_{\binom{m}{k}-1}$. If $\binom{m}{k} - 1 > m^2$, its image is nowhere dense in the probability simplex. This turns the common BSBM parameter-counting obstruction into a rigorous non-universality statement and connects the result directly to quantum generative modeling.

Unbounded-depth obstruction

Fixed-depth nowhere denseness and semi-algebraicity describe the topology and geometry of reachable coefficients of a model. However, they do not by themselves preclude the union over all depths from being dense. We therefore prove a separate topological obstruction.

Theorem 4 (Full-depth obstruction). *Assume $N \geq 2$. Define the full raw model by $\mathcal{H}_{\text{full}} := \bigcup_{L \geq 0} \mathcal{H}_L$. Then $\mathcal{H}_{\text{full}}$ is not dense in the $[0, 1]$ -valued center-invariant continuous target family*

$$C(U(N), \mathbb{R})_{[0,1]}^Z := \{f \in C(U(N), \mathbb{R})^Z : 0 \leq f \leq 1\}$$

with respect to the uniform norm. Moreover, it is not dense in the $[0, 1]$ -valued center-invariant polynomial target family $R(U(N), \mathbb{R})_{[0,1]}^Z$ either.

Proof sketch. Path connectedness of $U(N)$ makes $V_{L,W}(U) = W_L U W_{L-1} U \cdots W_1 U W_0$ homotopic to the power map $p_L(U) = U^L$. Since p_L has degree $L^N \neq 0$ on the closed oriented manifold $U(N)$, $V_{L,W}$ is surjective: for every fixed W and every $Q \in U(N)$, some $U \in U(N)$ satisfies $V_{L,W}(U) = Q$. Thus the readout is range-pinned to the full range $[0, 1]$. Uniform limits preserve minima and maxima, so positive-depth limits remain pinned to minimum 0 and maximum 1. Hence targets whose range stays strictly inside $(0, 1)$ are precluded. $\square$

This depth-independent obstruction rules out full universality of the raw single-Born-probability architecture. Overall, non-Abelian unitary-input re-uploading is governed by a genuine polynomial hierarchy: fixed depth is non-generic in each finite sector, and unbounded depth is not uniformly universal over the physically meaningful target space. The same non-representability and non-universality conclusions persist even for bounded-degree center-invariant polynomial targets of arbitrary degree.

References

- [1] A. Pérez-Salinas, A. Cervera-Lierta, E. Gil-Fuster, and J. I. Latorre, “Data re-uploading for a universal quantum classifier,” *Quantum* **4**, 226 (2020).

- [2] M. Schuld, R. Sweke, and J. J. Meyer, “The effect of data encoding on the expressive power of variational quantum-machine-learning models,” *Physical Review A* **103**, 032430 (2021).
- [3] M. C. Caro, E. Gil-Fuster, J. J. Meyer, J. Eisert, and R. Sweke, “Encoding-dependent generalization bounds for parametrized quantum circuits,” *Quantum* **5**, 582 (2021).
- [4] A. Kurkin et al., “Universality of classically trainable, quantum-deployed Boson Sampling Born Machines,” arXiv:2603.11014 (2026).
- [5] S. Aaronson and A. Arkhipov, “The computational complexity of linear optics,” *Theory of Computing* **9**, 143–252 (2013).
- [6] A. Pérez-Salinas, M. Yaghubi Rad, A. Barthe, and V. Dunjko, “Universal approximation of continuous functions with minimal quantum circuits,” *Physical Review Research* **7**, 043282 (2025).