

Re-uploading quantum data: a universal function approximator for quantum inputs

Hyunho Cha¹, Daniel K. Park^{2,3,4}, and Jungwoo Lee¹

¹NextQuantum and Department of Electrical and Computer Engineering, Seoul National University, Seoul 08826, Republic of Korea

²Department of Statistics and Data Science, Yonsei University, Seoul 03722, Republic of Korea

³Department of Applied Statistics, Yonsei University, Seoul 03722, Republic of Korea

⁴Department of Quantum Information, Yonsei University, Incheon 21983, Republic of Korea

Quantum data re-uploading has proved powerful for classical inputs, where repeatedly encoding features into a small circuit yields universal function approximation. Extending this idea to quantum inputs remains underexplored, as the information contained in a quantum state is not directly accessible in classical form. We propose and analyze a quantum data re-uploading architecture in which a qubit interacts sequentially with fresh copies of an arbitrary input state. The circuit can approximate any bounded continuous function using only one ancilla qubit and single-qubit measurements, provided that sufficiently many copies of the input state and, in general, sufficiently many layers are available. By alternating entangling unitaries with mid-circuit resets of the input register, the architecture realizes a discrete cascade of completely positive and trace-preserving maps, analogous to collision models in open quantum system dynamics. Our framework provides a space-efficient and expressive approach to designing quantum machine learning models that operate directly on quantum data.

1 Introduction

Quantum machine learning (QML) seeks to harness quantum computation to enhance machine learning tasks [1, 2, 3, 4, 5]. Quantum computers can perform certain linear algebra subroutines faster than classical machines under state preparation assumptions [6, 7, 8]. Motivated by such potential quantum speedups, a variety of QML models have been explored to outperform their classical counterparts, including quantum kernel methods and variational quantum circuits [9, 10, 11, 12, 13, 14, 15, 16].

A key component of any QML model is how data are encoded into and processed by quantum circuits [17, 18, 19, 20, 21, 22, 23, 24, 25, 26]. For classical input data, one common approach is to embed the data into a quantum state through parameterized gate operations. Recent work has shown that repeatedly encoding data within a circuit, known as data re-uploading, enhances a model's expressive power and can even make a single

Hyunho Cha: ovalavo@snu.ac.kr

Daniel K. Park: dkd.park@yonsei.ac.kr

Jungwoo Lee: junglee@snu.ac.kr

qubit a universal quantum classifier [27, 28, 18, 29]. The key idea is to interleave data-encoding gates with trainable quantum gates across several layers, effectively re-uploading the input at each layer. Similarly, quantum signal processing techniques were employed to prove that a sequence of single-qubit rotations interleaved with data-dependent operations can implement polynomial transformations of a classical input [30, 31, 32, 33].

Such results underscore that parameterizing quantum gates with classical data is a powerful paradigm in QML, enabling even small circuits to represent complex functions of the input data. Nonetheless, recent work has questioned whether quantum models trained on purely classical datasets can demonstrate genuine quantum advantage [34]. On the other hand, quantum advantage is more naturally expected when models are applied to quantum data [35] with concrete demonstrations emerging in recent studies [36]. However, when the input data are quantum rather than classical, this direct re-uploading approach encounters a fundamental obstacle. In a quantum input setting (where data are provided as quantum states), one cannot simply plug the input state’s “values” into rotation angles or other gate parameters, because those values are not explicitly available as classical numbers. Thus, a new strategy is required to re-upload quantum data in a multi-layer quantum model.

A closely related line of work studies direct multicopy estimation of nonlinear functionals of quantum states. Controlled-SWAP and cyclic-shift interferometric networks can estimate overlaps, purities, and spectral moments such as $\text{tr}(\rho^k)$ without first reconstructing ρ [37]. More generally, any m th-degree polynomial in the density matrix entries can be represented as the expectation value of an observable on m copies of ρ [38]. Recent quantum state function methods make this viewpoint sample-efficient by combining generalized SWAP tests with linear combinations of unitaries and parameterized circuits, enabling normalized degree- p polynomial functionals to be estimated to precision ε using $O(p/\varepsilon^2)$ copies, with applications to entropy, fidelity, and eigenvalue estimation [39]. Our goal is complementary: rather than designing a dedicated multicopy measurement for a specified functional, we analyze a general trainable sequential re-uploading architecture in which a single qubit coherently interacts with fresh copies of the input, yielding an ansatz with universal approximation guarantees for continuous functions of quantum states.

More specifically, we introduce a *quantum data re-uploading* architecture. The proposed model consists of a single-qubit *signal* register that interacts with the quantum input state in a sequence of layers. Instead of encoding quantum data as a gate parameter, each layer involves entangling the signal qubit with the input state, thereby uploading the input’s information into the joint system. In essence, the signal qubit acts as an intermediary that undergoes quantum operations that depend on both the input and learned parameters. By repeating this step across multiple layers, the circuit incrementally builds a complex transformation of the input state.

We show that this re-uploading architecture is capable of universal function approximation on quantum data while using only a single qubit for the processing unit. In fact, no matter how many qubits the input state contains, the model requires just one extra qubit (the signal register) to achieve universality. This result extends the spirit of single-qubit universality to the domain of quantum inputs.

A key feature of the proposed approach is its efficient use of qubits through qubit reuse [40, 41, 42, 43]. After the signal qubit interacts with the input in each layer, we simply discard the second register and then reset it to the input state. The signal register is left unmeasured so that it can continue to the subsequent layer. This qubit-reuse tactic allows the same register to load data in all layers, rather than needing a fresh input register for each layer. As a result, the total number of qubits required by the model does not scale with the number of layers, since only the input register and one working qubit are needed.




Figure 1: A single-qubit classical data re-uploading scheme. $U^{(l)}(\mathbf{x})$ denotes an input-dependent data encoding gate, while $V^{(l)}(\boldsymbol{\theta})$ denotes a trainable input-independent gate.

This approach is compatible with the capabilities of many modern quantum devices that support reset operations.

2 Notations

Let n be the number of qubits in the input system, and set $d = 2^n$. We denote the associated d -dimensional Hilbert space by $\mathcal{H}_d$. The set of all density operators on $\mathcal{H}_d$ is denoted as

$$\mathcal{D}(\mathcal{H}_d) = \{\rho \in \mathcal{B}(\mathcal{H}_d) \mid \rho = \rho^\dagger, \rho \succeq 0, \text{tr}(\rho) = 1\},$$

where $\mathcal{B}(\mathcal{H}_d)$ denotes the space of all bounded linear operators on $\mathcal{H}_d$. The corresponding space of observables is the vector space of Hermitian operators, which we denote by

$$\text{Herm}(\mathcal{H}_d) = \{O \in \mathcal{B}(\mathcal{H}_d) \mid O = O^\dagger\}.$$

The identity operator on $\mathcal{H}_d$ is denoted by $\mathbb{I}_d$. We adopt the notation

$$(\sigma^{(0)}, \sigma^{(1)}, \sigma^{(2)}, \sigma^{(3)}) = (\mathbb{I}_2, X, Y, Z)$$

to refer to the standard Pauli matrices.

3 Main results

Re-uploading classical input data into quantum circuits for function approximation is a well-established technique [27, 31, 28]. In this approach, classical input data $\mathbf{x}$ is encoded into a quantum circuit multiple times through parameterized gates, which alternate with trainable gates parameterized by $\boldsymbol{\theta}$ (see Figure 1). However, when the input data is quantum, its density matrix elements cannot be used directly to parameterize the circuit, as a full classical description of the state is generally not efficiently accessible. This naturally leads to the case of directly re-uploading the quantum data.

Let $\rho \in \mathcal{D}(\mathcal{H}_d)$ be input quantum data. Let A and B denote the signal and input registers, respectively, where ρ is loaded into B . The circuit illustrated in Figure 2a effectively accumulates information about ρ into A , which is initialized in the state $|0\rangle$. After L re-uploading layers, we obtain a final density matrix on register A . A classical post-processing step is then applied to extract a single real scalar output.

Remark 1. When $L = 1$, the re-uploading model resembles the Hadamard test. The circuit can be used to estimate the value of $\text{tr}(\rho U)$ for a given unitary $U \in \text{SU}(d)$ (see Figure 2b). For the maximally mixed input $\rho = \mathbb{I}_d/d$, the estimated quantity reduces to the normalized trace $\text{tr}(U)/d$, which appears in the trace estimation setting of the one clean qubit model of computation, also known as deterministic quantum computation with one clean qubit (DQC1) [44, 45]. This observation provides an informal connection between the $L = 1$ re-uploading circuit and DQC1 trace estimation.




Figure 2: (a) A general quantum data re-uploading scheme for universal function approximation. $\rho \in \mathcal{D}(\mathcal{H}_d)$ denotes the quantum input and gray boxes indicate arbitrary $SU(2d)$ elements. (b) The Hadamard test circuit can be used to estimate $\text{tr}(\rho U)$ by measuring the expectation value $\langle Z \rangle$ on system A, using two separate configurations: $V = H$ to obtain $\text{Re}(\text{tr}(\rho U))$, and $V = S^\dagger H$ to obtain $\text{Im}(\text{tr}(\rho U))$.



Figure 3: A restricted single-qubit data re-uploading scheme for universal function approximation. $\rho \in \mathcal{D}(\mathcal{H}_2)$ denotes the single-qubit input and gray boxes indicate arbitrary $SU(2)$ elements.

Remark 2. A standard variational quantum classifier (VQC) that processes only one copy of the quantum input state is contained in the $L = 1$ case of our quantum data re-uploading model, at least for the usual setting of a variational unitary followed by a final observable on the first qubit. Consequently, in tasks where the target function requires higher-order dependence on the input state and is achieved only for $L > 1$, the re-uploading architecture can represent functions that standard single-copy VQCs of this form cannot.

To realize the circuit shown in Figure 2a, we assume the ability to reset system B to the state ρ . Although qubit reuse remains relatively underexplored in quantum circuit optimization, this capability is gaining increasing attention as it can significantly reduce the total number of qubits required in a computation [41, 42]. As discussed in Section 4, the circuit in Figure 2a is logically equivalent to an $(Ln + 1)$ -qubit circuit, where L is the number of re-uploading layers. However, with qubit reuse, the same computation can be implemented using only $(n + 1)$ qubits, independent of L .

The reset of register B to ρ should be understood as fresh-copy access to the input state, rather than a standard hardware reset to a fixed reference state such as $|0\rangle^{\otimes n}$. Consequently, an L -layer circuit requires L preparations or copies of ρ . The total execution cost therefore depends on the state-preparation channel.

3.1 Re-uploading a single-parameter single-qubit state

We aim to establish that the re-uploading structure in Figure 2a possesses universal function approximation capability. Consider the simplest case where ρ is a single-parameter single-qubit pure state with real amplitudes:

$$\rho(t) = |\psi(t)\rangle\langle\psi(t)| \in \mathcal{D}(\mathcal{H}_2),$$

where

$$|\psi(t)\rangle = t|0\rangle + \sqrt{1-t^2}|1\rangle \in \mathcal{H}_2, \quad 0 < t < 1. \quad (1)$$

Also, let $\tau^{(l)}$ denote the density matrix of A after l re-uploading layers. Without loss of generality, we take $\tau^{(0)} = |+\rangle\langle+|$ in place of $|0\rangle\langle 0|$. We further simplify the scheme by restricting each $\text{SU}(4)$ element (a gray box in Figure 2a) to consist of a single-qubit gate on A followed by $\text{CNOT}_{B \rightarrow A}$ (see Figure 3). In this restricted form, the circuit structure mirrors the pattern of interleaving trainable gate layers with fixed, data-encoding operations, which in turn simplifies our subsequent analysis.

Let

$$\mathbf{r}^{(l)} = (r_1^{(l)}, r_2^{(l)}, r_3^{(l)}) = (\langle \sigma^{(1)} \rangle_{\tau^{(l)}}, \langle \sigma^{(2)} \rangle_{\tau^{(l)}}, \langle \sigma^{(3)} \rangle_{\tau^{(l)}}) \in \mathbb{R}^3$$

denote the Bloch vector of A after the l th re-uploading layer. Initially,

$$\tau^{(0)} = |+\rangle\langle+| \iff \mathbf{r}^{(0)} = (r_1^{(0)}, r_2^{(0)}, r_3^{(0)}) = (1, 0, 0).$$

At the l th re-uploading layer, the l th single-qubit gate $U^{(l)} \in \text{SU}(2)$ on A rotates the Bloch vector $\mathbf{r}^{(l-1)}$ by $\mathfrak{R}^{(l)} \in \text{SO}(3)$:

$$\mathbf{r}^{(l-1)} \mapsto \tilde{\mathbf{r}}^{(l)} = (\tilde{r}_1^{(l)}, \tilde{r}_2^{(l)}, \tilde{r}_3^{(l)}) = \mathfrak{R}^{(l)} \mathbf{r}^{(l-1)}.$$

The state corresponding to $\tilde{\mathbf{r}}^{(l)}$ is

$$\tilde{\tau}^{(l)} = \frac{1}{2} (\mathbb{I}_2 + \tilde{r}_1^{(l)} X + \tilde{r}_2^{(l)} Y + \tilde{r}_3^{(l)} Z).$$

Also, note that

$$\rho(t) = \begin{pmatrix} t^2 & t\sqrt{1-t^2} \\ t\sqrt{1-t^2} & 1-t^2 \end{pmatrix} = \frac{1}{2} (\mathbb{I}_2 + 2t\sqrt{1-t^2}X + (2t^2-1)Z)$$

and

$$\text{CNOT}_{B \rightarrow A} = (\mathbb{I}_2)_A \otimes |0\rangle\langle 0|_B + X_A \otimes |1\rangle\langle 1|_B.$$

By taking the tensor product $\tilde{\tau}^{(l)} \otimes \rho(t)$, applying $\text{CNOT}_{B \rightarrow A}$, and tracing out system B , the state $\tilde{\tau}^{(l)}$ evolves to

$$\begin{aligned} \tau^{(l)} &= ((\mathbb{I}_2)_A \otimes \langle 0|_B) \text{CNOT}_{B \rightarrow A} (\tilde{\tau}^{(l)} \otimes \rho(t)) \text{CNOT}_{B \rightarrow A} ((\mathbb{I}_2)_A \otimes |0\rangle_B) \\ &\quad + ((\mathbb{I}_2)_A \otimes \langle 1|_B) \text{CNOT}_{B \rightarrow A} (\tilde{\tau}^{(l)} \otimes \rho(t)) \text{CNOT}_{B \rightarrow A} ((\mathbb{I}_2)_A \otimes |1\rangle_B) \\ &= K_0 (\tilde{\tau}^{(l)} \otimes \rho(t)) K_0^\dagger + K_1 (\tilde{\tau}^{(l)} \otimes \rho(t)) K_1^\dagger, \end{aligned} \quad (2)$$

where the Kraus operators K_0 and K_1 are

$$\begin{aligned} K_0 &= ((\mathbb{I}_2)_A \otimes \langle 0|_B) \text{CNOT}_{B \rightarrow A} \\ &= ((\mathbb{I}_2)_A \otimes \langle 0|_B) ((\mathbb{I}_2)_A \otimes |0\rangle\langle 0|_B + X_A \otimes |1\rangle\langle 1|_B) \\ &= (\mathbb{I}_2)_A \otimes \langle 0|_B \end{aligned}$$

and

$$\begin{aligned} K_1 &= ((\mathbb{I}_2)_A \otimes \langle 1|_B) \text{CNOT}_{B \rightarrow A} \\ &= ((\mathbb{I}_2)_A \otimes \langle 1|_B) ((\mathbb{I}_2)_A \otimes |0\rangle\langle 0|_B + X_A \otimes |1\rangle\langle 1|_B) \\ &= X_A \otimes \langle 1|_B. \end{aligned}$$

After some calculations, we find

$$\begin{aligned} K_0 \left(\tilde{\tau}^{(l)} \otimes \rho(t) \right) K_0^\dagger &= \frac{1}{4} \left(\mathbb{I}_2 + \tilde{r}_1^{(l)} X + \tilde{r}_2^{(l)} Y + \tilde{r}_3^{(l)} Z \right) (1 + 2t^2 - 1) \\ &= \frac{1}{2} t^2 \left(\mathbb{I}_2 + \tilde{r}_1^{(l)} X + \tilde{r}_2^{(l)} Y + \tilde{r}_3^{(l)} Z \right) \end{aligned} \quad (3)$$

and

$$\begin{aligned} K_1 \left(\tilde{\tau}^{(l)} \otimes \rho(t) \right) K_1^\dagger &= \frac{1}{4} \left(\mathbb{I}_2 + \tilde{r}_1^{(l)} X - \tilde{r}_2^{(l)} Y - \tilde{r}_3^{(l)} Z \right) (1 - 2t^2 + 1) \\ &= \frac{1}{2} (1 - t^2) \left(\mathbb{I}_2 + \tilde{r}_1^{(l)} X - \tilde{r}_2^{(l)} Y - \tilde{r}_3^{(l)} Z \right). \end{aligned} \quad (4)$$

Let $\lambda = 2t^2 - 1 \in (-1, 1)$. Then from Eqs. (2), (3), and (4),

$$\tau^{(l)} = \frac{1}{2} \left(\mathbb{I}_2 + \tilde{r}_1^{(l)} X + \lambda \tilde{r}_2^{(l)} Y + \lambda \tilde{r}_3^{(l)} Z \right)$$

and therefore

$$\mathbf{r}^{(l)} = \left(\tilde{r}_1^{(l)}, \lambda \tilde{r}_2^{(l)}, \lambda \tilde{r}_3^{(l)} \right) = \tilde{\mathbf{r}}^{(l)} \odot (1, \lambda, \lambda),$$

where $\odot$ denotes the element-wise (Hadamard) product. That is, the y and z components of $\tilde{\mathbf{r}}^{(l)}$ are scaled by a factor of λ . In summary, the Bloch vector initially at $\mathbf{r}^{(0)} = (1, 0, 0)$ undergoes a sequence of alternating rotations and scalings.

Remark 3. In general, the Bloch vector undergoes a sequence of more complex transformations. Specifically, let $U \in \text{SU}(4)$. Then the evolution

$$\tau^{(l)} \mapsto \tau^{(l+1)} = \text{tr}_B \left(U \left(\tau^{(l)} \otimes \rho \right) U^\dagger \right)$$

corresponds to

$$\mathbf{r}^{(l)} \mapsto \mathbf{r}^{(l+1)} = M \mathbf{r}^{(l)} + \mathbf{d},$$

where $M \in \mathbb{R}^{3 \times 3}$ and $\mathbf{d} \in \mathbb{R}^3$ are given by

$$\begin{aligned} M_{ij} &= \frac{1}{2} \text{tr} \left[\left(\sigma^{(i)} \otimes \mathbb{I}_2 \right) U \left(\sigma^{(j)} \otimes \rho \right) U^\dagger \right], \\ d_i &= \frac{1}{2} \text{tr} \left[\left(\sigma^{(i)} \otimes \mathbb{I}_2 \right) U \left(\mathbb{I}_2 \otimes \rho \right) U^\dagger \right]. \end{aligned}$$

One can verify that this transformation reduces to a rotation followed by a scaling when the circuit is restricted to Figure 3.

Each $U^{(l)}$ can be further restricted to the form

$$R_z(\theta_l) = \begin{pmatrix} e^{-i\theta_l/2} & 0 \\ 0 & e^{i\theta_l/2} \end{pmatrix},$$

yielding the corresponding Bloch vector rotation

$$\mathfrak{R}^{(l)} = \begin{pmatrix} \cos \theta_l & -\sin \theta_l & 0 \\ \sin \theta_l & \cos \theta_l & 0 \\ 0 & 0 & 1 \end{pmatrix}. \quad (5)$$

Finally, we introduce learnable parameters $\mathbf{w} \in \mathbb{R}^3$ and $b \in \mathbb{R}$, so that the full set of parameters is given by $\boldsymbol{\theta} = (\theta_1, \dots, \theta_L, \mathbf{w}, b) \in \mathbb{R}^{L+4}$. The final output is computed via the mapping

$$\mathbf{r}^{(L)} \mapsto f_{\boldsymbol{\theta}}(\lambda) = \mathbf{w} \cdot \mathbf{r}^{(L)} + b \in \mathbb{R}, \quad (6)$$

where $\mathbf{r}^{(L)}$ depends on both $\boldsymbol{\theta}$ and λ . Note that $f_{\boldsymbol{\theta}}(\lambda)$ can be expressed as

$$\langle W \rangle_{\tau^{(L)}} + b = \text{tr}(\tau^{(L)} W) + b,$$

where

$$W = \sum_{i=1}^3 w_i \sigma^{(i)}$$

is a Hermitian observable. Thus, $f_{\boldsymbol{\theta}}(\lambda)$ can be evaluated by computing the expectation value of W with respect to the state $\tau^{(L)}$. The expectation values $\langle \sigma^{(1)} \rangle$ and $\langle \sigma^{(2)} \rangle$ can be accessed using only the ability to perform computational-basis measurements combined with single-qubit Clifford transformations. Moreover, the bias term b is introduced explicitly for convenience in subsequent analysis, but it can be absorbed into the observable as $W' = W + b \cdot \mathbb{I}_2$.

We can show that the function defined in Eq. (6) is capable of representing an arbitrary polynomial in λ , given a sufficient number of re-uploading layers L . Consider the following special cases:

$$\begin{aligned} \boldsymbol{\theta} &= (0, \dots, 0, \theta_l, 0, \dots, 0, \mathbf{e}_2, 0) \\ &\rightarrow \mathbf{r}^{(L)} = (\cos \theta_l, \lambda^{L+1-l} \sin \theta_l, 0), & f_{\boldsymbol{\theta}}(\lambda) &= \lambda^{L+1-l} \sin \theta_l, \\ \boldsymbol{\theta} &= (0, \dots, 0, \mathbf{e}_2, b) \\ &\rightarrow \mathbf{r}^{(L)} = (1, 0, 0), & f_{\boldsymbol{\theta}}(\lambda) &= b, \end{aligned} \quad (7)$$

where $\mathbf{e}_2 = (0, 1, 0)$. Since $f_{\boldsymbol{\theta}}(\lambda)$ is a polynomial of degree at most L , it can be represented by a coefficient vector $\mathbf{v}_{\boldsymbol{\theta}} \in \mathbb{R}^{L+1}$ such that

$$f_{\boldsymbol{\theta}}(\lambda) = \sum_{i=1}^{L+1} v_{\boldsymbol{\theta},i} \lambda^{i-1}.$$

Meanwhile, the Jacobian of $\mathbf{v}_{\boldsymbol{\theta}}$ at $\boldsymbol{\theta} = \boldsymbol{\theta}_0 = (0, \dots, 0, \mathbf{e}_2, 0)$ is

$$\mathbf{J}_{\mathbf{v}_{\boldsymbol{\theta}}}(\boldsymbol{\theta}_0) = \begin{bmatrix} 0 & 0 & \cdots & 0 & \boxed{\begin{matrix} 1 & 0 & 0 \\ 0 & 0 & 0 \end{matrix}} & 1 \\ 0 & 0 & \cdots & 1 & \boxed{\begin{matrix} 0 & 0 & 0 \\ 0 & 0 & 0 \end{matrix}} & 0 \\ \vdots & \vdots & \ddots & \vdots & \vdots & \vdots \\ 0 & 1 & \cdots & 0 & \boxed{\begin{matrix} 0 & 0 & 0 \\ 0 & 0 & 0 \end{matrix}} & 0 \\ 1 & 0 & \cdots & 0 & \boxed{\begin{matrix} 0 & 0 & 0 \\ 0 & 0 & 0 \end{matrix}} & 0 \end{bmatrix} \in \mathbb{R}^{(L+1) \times (L+4)}, \quad (8)$$

where the boxed submatrix represents the Jacobian with respect to $\mathbf{w}$. It is evident that $\mathbf{J}_{\mathbf{v}_{\boldsymbol{\theta}}}(\boldsymbol{\theta}_0)$ has full rank. If we fix $\mathbf{w} = \mathbf{e}_2$ and vary only the remaining $L+1$ parameters $\boldsymbol{\theta}' = (\theta_1, \dots, \theta_L, b)$, we obtain the reduced map

$$\mathbf{v}'_{\boldsymbol{\theta}'} := \mathbf{v}_{(\theta_1, \dots, \theta_L, \mathbf{e}_2, b)}.$$

The Jacobian of $\mathbf{v}'_{\boldsymbol{\theta}'}$ at $\boldsymbol{\theta}' = \mathbf{0}$ is the matrix obtained by removing the boxed submatrix from $\mathbf{J}_{\mathbf{v}_{\boldsymbol{\theta}}}(\boldsymbol{\theta}_0)$ in Eq. (8), which is invertible.

Theorem 1 ([46], Theorem 1.1.7). *If f is a continuously differentiable function from an open subset $\tilde{S} \subset \mathbb{R}^n$ into $\mathbb{R}^n$, and the derivative $f'(p)$ is invertible at a point p , then there exist neighborhoods S of p in $\tilde{S}$ and T of $q = f(p)$ such that $f(S) \subset T$ and $f : S \rightarrow T$ is bijective.*

Note that $\mathbf{v}'_0 = \mathbf{0}$. Then by Theorem 1, for arbitrary $\mathbf{v} \in \mathbb{R}^{L+1}$, there exists a real constant $\varepsilon \neq 0$ and a point $\boldsymbol{\theta}' \in \mathbb{R}^{L+1}$ near $\mathbf{0}$ such that

$$\mathbf{v}'_{\boldsymbol{\theta}'} = \varepsilon \mathbf{v}.$$

However, we can scale $\mathbf{v}'_{\boldsymbol{\theta}'}$ by an arbitrary constant, since

$$f_{(\theta_1, \dots, \theta_L, c\mathbf{w}, cb)}(\lambda) = cf_{(\theta_1, \dots, \theta_L, \mathbf{w}, b)}(\lambda). \quad (9)$$

Choosing $c = \varepsilon^{-1}$, Eq. (9) implies that there exists a point $\boldsymbol{\theta} \in \mathbb{R}^{L+1}$ such that $\mathbf{v}_{\boldsymbol{\theta}} = \mathbf{v}$.

3.2 Re-uploading a single-parameter multi-qubit state

This section generalizes the single-parameter construction in Section 3.1 from a single-qubit input to a multi-qubit input. Any n -qubit density matrix can be represented as

$$\rho = \frac{1}{d} \left(\mathbb{I}_d + \sum_{\alpha=1}^{d^2-1} \lambda_{\alpha} W_{\alpha} \right), \quad (10)$$

where W_{α} are generalized Pauli matrices (n -fold tensor products of $\sigma^{(0)}$, $\sigma^{(1)}$, $\sigma^{(2)}$, and $\sigma^{(3)}$). Let W_{α}^{+} and W_{α}^{-} denote the projectors onto the positive and negative eigenspaces of W_{α} . Because every generalized Pauli is a traceless Hermitian unitary, its only eigenvalues ± 1 must occur equally often, so the $+1$ - and -1 -eigenspaces each span half of $\mathcal{H}_d$. Therefore, W_{α}^{+} and W_{α}^{-} can be expressed as

$$W_{\alpha}^{+} = \frac{1}{2}(\mathbb{I}_d + W_{\alpha}) = \sum_{\beta=1}^{d/2} |\psi_{\beta}^{(\alpha)}\rangle \langle \psi_{\beta}^{(\alpha)}|$$

and

$$W_{\alpha}^{-} = \frac{1}{2}(\mathbb{I}_d - W_{\alpha}) = \sum_{\gamma=1}^{d/2} |\phi_{\gamma}^{(\alpha)}\rangle \langle \phi_{\gamma}^{(\alpha)}|,$$

where $\{|\psi_{\beta}^{(\alpha)}\rangle\}_{\beta}$ and $\{|\phi_{\gamma}^{(\alpha)}\rangle\}_{\gamma}$ are orthonormal families of states. Define the multi-qubit controlled unitary

$$\text{CU}_{B \rightarrow A}^{(\alpha)} = (\mathbb{I}_2)_A \otimes (W_{\alpha}^{+})_B + X_A \otimes (W_{\alpha}^{-})_B$$

and the Kraus operators

$$\begin{aligned} K_{\beta}^{(\alpha+)} &= \left((\mathbb{I}_2)_A \otimes \langle \psi_{\beta}^{(\alpha)} |_B \right) \text{CU}_{B \rightarrow A}^{(\alpha)} = (\mathbb{I}_2)_A \otimes \langle \psi_{\beta}^{(\alpha)} |_B, \\ K_{\gamma}^{(\alpha-)} &= \left((\mathbb{I}_2)_A \otimes \langle \phi_{\gamma}^{(\alpha)} |_B \right) \text{CU}_{B \rightarrow A}^{(\alpha)} = X_A \otimes \langle \phi_{\gamma}^{(\alpha)} |_B. \end{aligned}$$

Then the state

$$\tilde{\tau}^{(l)} \otimes \rho = \frac{1}{2d} \left(\mathbb{I}_2 + \tilde{r}_1^{(l)} X + \tilde{r}_2^{(l)} Y + \tilde{r}_3^{(l)} Z \right) \otimes \left(\mathbb{I}_d + \sum_{\alpha'=1}^{d^2-1} \lambda_{\alpha'} W_{\alpha'} \right)$$

evolves to

$$\tau^{(l)} = \frac{1}{2} \left(\mathbb{I}_2 + \tilde{r}_1^{(l)} X + [\lambda_{\alpha} \tilde{r}_2^{(l)}] Y + [\lambda_{\alpha} \tilde{r}_3^{(l)}] Z \right), \quad (11)$$

where the subsystem labels A and B have been omitted for notational simplicity (see Appendix A for the derivation). That is, the y and z components of $\tilde{\mathbf{r}}^{(l)}$ are scaled by

a factor of λ_α . If the target function of ρ depends only on a single parameter $\lambda_{\alpha'}$, then the proof of universality can be completed by employing the same approach outlined in Section 3.1 with a target polynomial

$$f(\lambda_{\alpha'}) = c_0 + \sum_{\omega=1}^{\Omega} c_\omega \lambda_{\alpha'}^\omega.$$

However, this argument is intrinsically single-parameter. For a fixed choice of α , the channel induced by $\text{CU}_{B \rightarrow A}^{(\alpha)}$ uses the input state ρ only through the single expectation value $\lambda_\alpha = \text{tr}(\rho W_\alpha)$. Consequently, the resulting one-qubit dynamics only produces expressions that can be organized as univariate polynomials in λ_α . It does not by itself establish universality for a general function $f(\lambda_1, \dots, \lambda_{d^2-1})$, because such a function may require arbitrary multivariate monomials, such as $\lambda_{\alpha_1} \lambda_{\alpha_2} \cdots \lambda_{\alpha_k}$, involving different generalized Pauli components. The above construction fails to show how different components can be accessed, stored, mixed, and multiplied within the same re-uploading architecture to generate a dense class of multivariate polynomials. Moreover, it also does not yield a uniform guarantee that the final readout observable can be chosen with controlled operator norm. The next section resolves these issues by proving that our architecture can simulate ordinary *classical* data re-uploading.

3.3 Re-uploading an arbitrary multi-qubit state

In this section, we use the re-uploading architecture of Figure 2a to simulate an ordinary single-qubit classical data re-uploading circuit whose classical input is λ .

We use the standard normalization

$$\text{tr}(W_\alpha W_\beta) = d\delta_{\alpha\beta}, \quad \text{tr}(W_\alpha) = 0.$$

Therefore the coefficients in Eq. (10) are recovered as expectation values:

$$\begin{aligned} \text{tr}(\rho W_\beta) &= \frac{1}{d} \text{tr} \left[\left(\mathbb{I}_d + \sum_{\alpha=1}^{d^2-1} \lambda_\alpha W_\alpha \right) W_\beta \right] \\ &= \frac{1}{d} \left(\text{tr}(W_\beta) + \sum_{\alpha=1}^{d^2-1} \lambda_\alpha \text{tr}(W_\alpha W_\beta) \right) \\ &= \lambda_\beta. \end{aligned} \tag{12}$$

Since W_α has eigenvalues in $\{+1, -1\}$, Eq. (12) also implies $|\lambda_\alpha| \leq 1$ for every physical input state ρ . Since each W_α is Hermitian and unitary, it is involutory:

$$W_\alpha^\dagger = W_\alpha, \quad W_\alpha^2 = \mathbb{I}_d.$$

Similarly, for each signal Pauli $\sigma^{(i)}$, $i \in \{1, 2, 3\}$,

$$(\sigma^{(i)})^\dagger = \sigma^{(i)}, \quad (\sigma^{(i)})^2 = \mathbb{I}_2.$$

Fix a generalized Bloch coordinate λ_α and a signal Pauli direction $i \in \{1, 2, 3\}$. For a real number δ , define the $(n+1)$ -qubit unitary

$$U_{\alpha,i}(\delta) = \exp\left(-i\delta \sigma_A^{(i)} \otimes (W_\alpha)_B\right). \tag{13}$$

This acts only on the signal register A and the current copy of ρ in register B . If

$$H_{\alpha,i} = \sigma_A^{(i)} \otimes (W_\alpha)_B,$$

then

$$\text{tr}(H_{\alpha,i}) = \text{tr}(\sigma^{(i)})\text{tr}(W_\alpha) = 0,$$

so $\det(U_{\alpha,i}(\delta)) = \exp(-i\delta\text{tr}(H_{\alpha,i})) = 1$. Thus $U_{\alpha,i}(\delta) \in \text{SU}(2d)$. Also,

$$H_{\alpha,i}^2 = (\sigma_A^{(i)})^2 \otimes ((W_\alpha)_B)^2 = (\mathbb{I}_2)_A \otimes (\mathbb{I}_d)_B.$$

Expanding the exponential in Eq. (13) gives

$$U_{\alpha,i}(\delta) = \cos \delta (\mathbb{I}_2)_A \otimes (\mathbb{I}_d)_B - i \sin \delta \sigma_A^{(i)} \otimes (W_\alpha)_B \quad (14)$$

(see Appendix B).

Let $\tau \in \mathcal{D}(\mathcal{H}_2)$ be the current state of the signal register. The re-uploading layer induced by Eq. (13) is the channel

$$\Phi_{\alpha,i}^{(\delta)}(\tau) = \text{tr}_B \left[U_{\alpha,i}(\delta) (\tau \otimes \rho) U_{\alpha,i}(\delta)^\dagger \right]. \quad (15)$$

For compactness, set

$$c_\delta = \cos \delta, \quad s_\delta = \sin \delta, \quad S_i = \sigma^{(i)}.$$

Using Eq. (14) and suppressing subsystem labels, we have

$$\begin{aligned} U_{\alpha,i}(\delta) (\tau \otimes \rho) U_{\alpha,i}(\delta)^\dagger &= (c_\delta \mathbb{I} - i s_\delta S_i \otimes W_\alpha) (\tau \otimes \rho) (c_\delta \mathbb{I} + i s_\delta S_i \otimes W_\alpha) \\ &= c_\delta^2 (\tau \otimes \rho) + s_\delta^2 (S_i \tau S_i \otimes W_\alpha \rho W_\alpha) \\ &\quad + i c_\delta s_\delta [(\tau S_i \otimes \rho W_\alpha) - (S_i \tau \otimes W_\alpha \rho)]. \end{aligned} \quad (16)$$

Now take the partial trace over register B . The first term gives

$$\text{tr}_B(\tau \otimes \rho) = \tau \text{tr}(\rho) = \tau.$$

The second term gives

$$\begin{aligned} \text{tr}_B(S_i \tau S_i \otimes W_\alpha \rho W_\alpha) &= S_i \tau S_i \text{tr}(W_\alpha \rho W_\alpha) \\ &= S_i \tau S_i \text{tr}(\rho W_\alpha^2) \\ &= S_i \tau S_i. \end{aligned}$$

For the two cross terms, Eq. (12) gives

$$\begin{aligned} \text{tr}_B(\tau S_i \otimes \rho W_\alpha) &= \tau S_i \text{tr}(\rho W_\alpha) = \lambda_\alpha \tau S_i, \\ \text{tr}_B(S_i \tau \otimes W_\alpha \rho) &= S_i \tau \text{tr}(W_\alpha \rho) = \lambda_\alpha S_i \tau. \end{aligned}$$

Substituting these identities into Eq. (16), we obtain

$$\begin{aligned} \Phi_{\alpha,i}^{(\delta)}(\tau) &= c_\delta^2 \tau + s_\delta^2 S_i \tau S_i + i c_\delta s_\delta \lambda_\alpha (\tau S_i - S_i \tau) \\ &= c_\delta^2 \tau + s_\delta^2 S_i \tau S_i - i c_\delta s_\delta \lambda_\alpha [S_i, \tau]. \end{aligned} \quad (17)$$

This identity shows that a single re-uploading layer produces, to first order in δ , a unitary rotation of the signal qubit whose angle is proportional to λ_α .

Indeed, Eq. (17) gives

$$\Phi_{\alpha,i}^{(\delta)}(\tau) = \tau - i\delta\lambda_\alpha[S_i, \tau] + \delta^2(S_i\tau S_i - \tau) + O(\delta^3). \quad (18)$$

The term $-i\delta\lambda_\alpha[S_i, \tau]$ is exactly the first-order action of the data-dependent rotation

$$R_{\alpha,i}(t; \boldsymbol{\lambda}) = \exp(-it\lambda_\alpha S_i). \quad (19)$$

Let

$$\mathcal{R}_{\alpha,i}^{(t)}(\tau) = R_{\alpha,i}(t; \boldsymbol{\lambda})\tau R_{\alpha,i}(t; \boldsymbol{\lambda})^\dagger. \quad (20)$$

Since $S_i^2 = \mathbb{I}_2$,

$$R_{\alpha,i}(t; \boldsymbol{\lambda}) = \cos(t\lambda_\alpha)\mathbb{I}_2 - i\sin(t\lambda_\alpha)S_i.$$

Hence the classical data-encoding rotation is

$$\mathcal{R}_{\alpha,i}^{(t)}(\tau) = \cos^2(t\lambda_\alpha)\tau + \sin^2(t\lambda_\alpha)S_i\tau S_i - i\sin(t\lambda_\alpha)\cos(t\lambda_\alpha)[S_i, \tau]. \quad (21)$$

Comparing Eq. (17) at $\delta = t$ with Eq. (21), the first-order terms agree. The difference begins at second order in the small interaction angle. More explicitly, for $|\lambda_\alpha| \leq 1$,

$$\begin{aligned} \Phi_{\alpha,i}^{(\delta)}(\tau) - \mathcal{R}_{\alpha,i}^{(\delta)}(\tau) &= \left(\cos^2 \delta - \cos^2(\delta\lambda_\alpha)\right)\tau + \left(\sin^2 \delta - \sin^2(\delta\lambda_\alpha)\right)S_i\tau S_i \\ &\quad - i(\lambda_\alpha \sin \delta \cos \delta - \sin(\delta\lambda_\alpha)\cos(\delta\lambda_\alpha))[S_i, \tau], \end{aligned} \quad (22)$$

and this difference is uniformly $O(\delta^2)$ over all input states ρ and all signal states τ : there exists a constant $C > 0$, independent of ρ , τ , α , and i , such that for all $|\delta| \leq 1$,

$$\left\|\Phi_{\alpha,i}^{(\delta)}(\tau) - \mathcal{R}_{\alpha,i}^{(\delta)}(\tau)\right\|_1 \leq C\delta^2. \quad (23)$$

A finite-angle classical data-dependent rotation is obtained from many re-uploading layers. Fix a desired rotation angle parameter $t \in \mathbb{R}$, choose an integer $N \geq 1$, and set $\delta = t/N$. The re-uploading architecture can apply the channel $\Phi_{\alpha,i}^{(t/N)}$ repeatedly N times, each time using a fresh copy of ρ in register B . Since the ideal channels satisfy

$$\left(\mathcal{R}_{\alpha,i}^{(t/N)}\right)^N = \mathcal{R}_{\alpha,i}^{(t)},$$

we have

$$\lim_{N \rightarrow \infty} \left(\Phi_{\alpha,i}^{(t/N)}\right)^N(\tau) = \mathcal{R}_{\alpha,i}^{(t)}(\tau)$$

uniformly in ρ and τ . Therefore, the quantum data re-uploading architecture can approximate the classical data-dependent single-qubit rotation

$$\tau \mapsto \exp(-it\lambda_\alpha S_i)\tau \exp(it\lambda_\alpha S_i)$$

by using many weak interactions.

We next embed this primitive into an ordinary single-qubit classical data re-uploading circuit. Consider a classical data re-uploading model whose classical input is

$$\boldsymbol{\lambda} = (\lambda_1, \dots, \lambda_{d^2-1}) \in \mathbb{R}^{d^2-1}.$$

An instance of such a model with finite depth has the form

$$\mathcal{C}_\theta(\boldsymbol{\lambda}) = \mathcal{V}^{(M)} \circ \mathcal{R}_{\alpha_M, i_M}^{(t_M)} \circ \mathcal{V}^{(M-1)} \circ \dots \circ \mathcal{R}_{\alpha_1, i_1}^{(t_1)} \circ \mathcal{V}^{(0)}, \quad (24)$$

where each $\mathcal{V}^{(m)}$ is an input-independent single-qubit unitary channel,

$$\mathcal{V}^{(m)}(\tau) = V^{(m)}\tau(V^{(m)})^\dagger, \quad V^{(m)} \in \text{SU}(2),$$

and each $\mathcal{R}_{\alpha_m, i_m}^{(t_m)}$ is defined as in Eq. (20). The symbol $\boldsymbol{\theta}$ denotes the continuous trainable parameters appearing in the circuit and readout. The Pauli directions i_m and the coordinates α_m are discrete architecture choices. This proves that every finite ordinary classical data re-uploading circuit on the generalized Bloch vector $\boldsymbol{\lambda}$ can be approximated by a single quantum data re-uploading circuit. Therefore, the quantum data re-uploading architecture inherits the universal approximation capability of ordinary single-qubit classical data re-uploading on the compact generalized Bloch-vector domain [47].

Theorem 2 (Universality of quantum data re-uploading). *Let $n \in \mathbb{N}$, let $d = 2^n$, and set $m = d^2 - 1$. For $\rho \in \mathcal{D}(\mathcal{H}_d)$, write*

$$\boldsymbol{\lambda}(\rho) = (\lambda_1(\rho), \dots, \lambda_m(\rho)), \quad \lambda_\alpha(\rho) = \text{tr}(\rho W_\alpha),$$

and let

$$\Lambda_d = \{\boldsymbol{\lambda}(\rho) : \rho \in \mathcal{D}(\mathcal{H}_d)\} \subset [-1, 1]^m.$$

Consider the following input access and circuit model.

1. The input is given by fresh copy access to an unknown state $\rho \in \mathcal{D}(\mathcal{H}_d)$. Equivalently, before each upload the input register $B \simeq \mathcal{H}_d$ can be reset to ρ . No classical description of ρ or of $\boldsymbol{\lambda}(\rho)$ is assumed.
2. The signal register is a single qubit $A \simeq \mathcal{H}_2$, initialized in $\tau^{(0)} = |0\rangle\langle 0|$. The allowed gates are arbitrary single-qubit unitaries $V \in \text{SU}(2)$ on A and the upload gates

$$U_{\alpha, i}(\delta) = \exp\left(-i\delta \sigma_A^{(i)} \otimes (W_\alpha)_B\right), \quad \alpha \in \{1, \dots, m\}, \quad i \in \{1, 2, 3\}, \quad \delta \in \mathbb{R}.$$

3. After each upload gate, register B is discarded or reset, while A is kept coherent.
4. Only the final signal qubit A is measured. The allowed readout is an affine combination of single-qubit Pauli expectations,

$$h(\rho) = b_{\text{out}} + \sum_{i=1}^3 a_i \text{tr}\left(\sigma^{(i)} \tau_A^{\text{out}}(\rho)\right) = b_{\text{out}} + \text{tr}\left(O_{\text{out}} \tau_A^{\text{out}}(\rho)\right),$$

$$\text{where } O_{\text{out}} = \sum_{i=1}^3 a_i \sigma^{(i)} \in \text{Herm}(\mathcal{H}_2).$$

Then the model is a universal approximator for continuous functions of ρ in the uniform norm. More precisely, for every bounded continuous function $F : \mathcal{D}(\mathcal{H}_d) \rightarrow \mathbb{R}$ and every $\varepsilon > 0$, there exist a finite number L_Q of uploads, choices of (α, i, δ) for those uploads, single-qubit unitaries on A , and a readout $(O_{\text{out}}, b_{\text{out}})$ such that

$$\sup_{\rho \in \mathcal{D}(\mathcal{H}_d)} \left| b_{\text{out}} + \text{tr}\left(O_{\text{out}} \tau_A^{\text{out}}(\rho)\right) - F(\rho) \right| < \varepsilon$$

and

$$\|O_{\text{out}}\|_\infty \leq \frac{1}{2} \left(\sup_{\rho \in \mathcal{D}(\mathcal{H}_d)} F(\rho) - \inf_{\rho \in \mathcal{D}(\mathcal{H}_d)} F(\rho) \right). \quad (25)$$




Figure 4: Two equivalent representations of our re-uploading model.

Furthermore, the approximation can be chosen by first selecting a finite single-qubit classical re-uploading circuit of the form given in Eq. (24), and then replacing each data-dependent rotation $\mathcal{R}_{\alpha_j, i_j}^{(t_j)}$ by N_j weak upload layers $(\Phi_{\alpha_j, i_j}^{(t_j/N_j)})^{N_j}$. If $N_j \geq |t_j|$ for all j , the simulated quantum channel

$$\hat{\mathcal{C}}_{\theta, \mathbf{N}}^{\rho} = \mathcal{V}^{(M)} \circ (\Phi_{\alpha_M, i_M}^{(t_M/N_M)})^{N_M} \circ \mathcal{V}^{(M-1)} \circ \dots \circ (\Phi_{\alpha_1, i_1}^{(t_1/N_1)})^{N_1} \circ \mathcal{V}^{(0)}$$

satisfies the uniform trace-norm simulation bound

$$\sup_{\rho \in \mathcal{D}(\mathcal{H}_d)} \sup_{\tau \in \mathcal{D}(\mathcal{H}_2)} \left\| \hat{\mathcal{C}}_{\theta, \mathbf{N}}^{\rho}(\tau) - \mathcal{C}_{\theta}(\lambda(\rho))(\tau) \right\|_1 \leq 6 \sum_{j=1}^M \frac{t_j^2}{N_j}. \quad (26)$$

Consequently, the total number of copies used is $L_Q = \sum_{j=1}^M N_j$.

Proof. See Appendix C. □

Remark 4. The result above should be interpreted as an existential width-efficient expressivity statement, not as an efficient $\text{SU}(2d)$ circuit-synthesis theorem. In other words, we do not prove that the number of upload layers, elementary gates, or sample complexity scale efficiently for arbitrary tasks. Thus the theorem establishes universality with small qubit width, whereas the design of explicit and depth-efficient circuit families for useful classes of functions remains a separate problem.

4 Interpretation as an entangling measurement

By the deferred measurement principle, the circuit in Figure 2a with L re-uploading layers is equivalent to an $(Ln + 1)$ -qubit circuit that prepares L copies of ρ , applies a sequence of L $(n + 1)$ -qubit unitaries, and performs measurements on the signal register (see Figure 4).

This reinterpretation provides a perspective closely related to Ref. [36]. In both cases, quantum memory enables entangling measurements across multiple copies of a state and provides capabilities far beyond sequential single-copy measurements with classical post-processing. Since the signal register serves as the quantum memory that carries information between successive uploads, we expect the potential advantage of the architecture to become more pronounced when this register is enlarged from a single qubit to a larger system. In our architecture (Figure 4b), the unitary transformation enclosed in the dashed

box can be absorbed into the measurement. This enables the approximation of arbitrary continuous functions of the target quantum input state (given sufficient copies and depth), a task that is more difficult when restricted to local measurements followed by classical post-processing.

To efficiently train the circuit in Figure 4b, one can apply parameter-shift rules, which express quantum derivatives via function evaluations at shifted inputs [48, 10]. Given the logical equivalence to the re-uploading model in Figure 4a, and our focus on observables in system A , these rules can be applied just as effectively within the re-uploading formulation. However, in our numerical simulations (Section 5), parameter-shift rules do not directly apply due to the general full-unitary parameterizations. We instead employ automatic differentiation to evaluate the gradients.

We briefly discuss the case of machine learning on a sequence of quantum states, i.e., on approximating arbitrary functions of

$$\mathcal{S} = (\rho^{(1)}, \dots, \rho^{(N)}).$$

It follows directly from the previous sections that re-uploading the tensor product state

$$\rho_{\text{prod}} = \rho^{(1)} \otimes \dots \otimes \rho^{(N)}$$

enables approximation of any bounded continuous function of the sequence $\mathcal{S}$. However, qubit resources can be used more efficiently. Instead of uploading the full tensor product at once, it suffices to upload a single $\rho^{(i)}$ at each layer while still preserving universality. To see this, let $\rho^{(i)} \in \mathcal{D}(\mathcal{H}_{d_i})$ with Bloch vector coefficients

$$\boldsymbol{\lambda}^{(i)} = \left(\lambda_{\alpha}^{(i)} \right)_{\alpha=1}^{d_i^2-1}.$$

The joint state ρ_{prod} is fully determined by the concatenated coefficient vector

$$\tilde{\boldsymbol{\lambda}} = \left\| \boldsymbol{\lambda}^{(i)} \right\|_{i=1}^N \in \mathbb{R}^{\sum_{i=1}^N d_i^2 - N}.$$

Consequently, any bounded continuous function of ρ_{prod} can be uniformly approximated by a polynomial in the variables of $\tilde{\boldsymbol{\lambda}}$. At each layer, we may select a variable $\lambda_{\alpha}^{(i)}$ and approximate the corresponding data-dependent rotation of the signal qubit. By iterating across layers and across different states $\rho^{(i)}$, one recovers the same polynomial approximation structure as in the single-state case, requiring at most $1 + \lceil \max_i \log_2 d_i \rceil$ qubits.

This perspective also resonates with the process learning setting discussed in Ref. [36], where access to quantum memory enables joint measurements across a sequence of states, providing a fundamental advantage over independent single-copy measurements.

5 Experiments

While the preceding sections establish universality through infinitesimal and existence arguments, such proofs do not imply uniqueness of the corresponding parameter solutions. In practice, training the re-uploading circuit may yield different parameter configurations that realize the same target function. To investigate this trainability aspect beyond pure existence proofs, we conduct a series of experiments in this section, including both regression and classification tasks.





Figure 5: Polynomial fitting results for the state in Eq. (1), where $\lambda = 2t^2 - 1$. Solid green lines denote the target functions, and dotted blue lines denote the fitted results obtained with our re-uploading model. (a) Results for training the restricted re-uploading circuit in Figure 3. The target function to be fitted is shown at the top of each plot. From left to right: plots with $L = 1$, $L = 4$, and $L = 2$ layers, respectively. (b) Approximation without training, using Eq. (27) for different values of Δ .

Implementation details. In all experiments, rather than using a specific circuit ansatz, we employ a full unitary parameterization via $U = \exp(iH)$, where $H \in \text{Herm}(\mathcal{H}_{2d})$. This full-unitary parameterization is adopted for numerical flexibility rather than hardware efficiency (see Remark 4). Training was implemented in PyTorch using full-batch optimization. We used the Adam optimizer with learning rate 10^{-2} . As a heuristic, all real parameters were initialized independently from $\mathcal{N}(0, 1)$. The signal qubit was initialized to $|0\rangle$ in all runs. For polynomial fitting, the regression grid consists of 201 evenly spaced points $\lambda_i \in [-1, 1]$. The regression loss is the mean squared error. The models were trained for 10^4 epochs. For all classification experiments, the training and test set sizes are $N_{\text{train}} = 1000$ and $N_{\text{test}} = 500$. The model is trained for 1000 epochs against binary labels $y_i \in \{0, 1\}$ using the mean squared error. Code is available in [49].

5.1 Polynomial fitting

As a simple example, we trained the restricted re-uploading circuit in Figure 3 using the input states in Eq. (1). The only trainable elements are sequences of single-qubit unitaries, and our targets are three univariate functions in the variable $\lambda = 2t^2 - 1$ (see Figure 5). The results are shown in Figure 5a.

We provide an alternative perspective on the re-uploading model's approximation capabilities. Let $\mathbf{v} \in \mathbb{R}^{L+1}$ and define

$$\boldsymbol{\theta}(\mathbf{v}; \Delta) = \left(v_{L+1}\Delta, \dots, v_2\Delta, \frac{\mathbf{e}_2}{\Delta}, v_1 \right) \in \mathbb{R}^{L+4}. \quad (27)$$

Then from Eq. (7), we see that

$$\lim_{\Delta \rightarrow 0} f_{\boldsymbol{\theta}(\mathbf{v}; \Delta)} = \sum_{i=1}^{L+1} v_i \lambda^{i-1}.$$

Note that the parameters in Eq. (27) are for illustrative purposes only, as the true polynomial coefficients $\mathbf{v}$ are unknown in practical scenarios. If $\mathbf{v}$ is known, then the parameters in Eq. (27) can be directly computed by fixing Δ , without requiring training. However, taking $\Delta \rightarrow 0$ is impractical, since the factor $1/\Delta$ diverges and consequently drives the weight vector $\mathbf{w}$ in Eq. (6) to blow up. This leads to poor numerical behavior, especially when only a limited number of measurement shots are available to estimate expectation values of observables. In practice, we observe that training the parameters naturally avoids this small- Δ regime, while still achieving comparable approximation performance to that obtained from Eq. (27) with extremely small values of Δ . Approximation results are shown in Figure 5b.

5.2 Purity classification

In this section, we consider the task of purity classification for single-qubit states. The purity of a density matrix ρ is defined by $\text{tr}(\rho^2)$, satisfying $1/d \leq \text{tr}(\rho^2) \leq 1$ [50]. We sample single-qubit states by choosing Bloch vectors uniformly from the unit ball in $\mathbb{R}^3$, and formulate a binary classification problem by asking whether the purity of ρ exceeds the threshold $\text{th} = (1 + 2^{-2/3})/2 \approx 0.815$, assigning class 1 if it does and class 0 otherwise. The value $(1 + 2^{-2/3})/2$ ensures balanced classes, i.e., $P(\text{tr}(\rho^2) < \text{th}) = P(\text{tr}(\rho^2) \geq \text{th}) = 0.5$. We generated a training set of 1000 states and a test set of 500 states, and trained the unrestricted re-uploading circuit in Figure 2a to perform the classification. Results are shown in Figures 6a and 7.




Figure 6: Test accuracy curves for (a) purity classification and (b) entanglement entropy classification with $L = 1, 2, 3, 4$.




Figure 7: Purity classification histograms for re-uploading depths $L = 1$ to $L = 4$, where the corresponding test accuracies are 0.51, 0.61, 0.95, and 0.98, respectively. The x -axis shows purity and the y -axis the sample count. Blue bars indicate predictions of class 1 and red bars predictions of class 0.

The re-uploading model is strictly less expressive than a global $(Ln + 1)$ -qubit unitary, even if we restrict our focus to system A . For example, we can implement the *purity test* circuit by allowing an entangling operation between just two copies of ρ (see Appendix D for details), which is a $(2n + 1)$ -qubit unitary. However, the same test cannot be implemented by the re-uploading model with $L = 2$ (which also uses two copies of ρ), as empirically demonstrated in Figure 7b and formally stated in the following observation.

Observation 1. *The re-uploading model with $L = 2$, where the output has the form of Eq. (6), cannot implement the purity function.*

Proof. See Appendix E. □

5.3 Entanglement entropy classification

One important nonlinear function of ρ is the second-order Rényi entanglement entropy

$$-\log(\text{tr}(\rho^{\otimes 2} S_A)),$$

where A is a subsystem of ρ and the local swap operator S_A acts as

$$S_A|i_{AJA^c}\rangle|k_{AJA^c}\rangle = |k_{AJA^c}\rangle|i_{AJA^c}\rangle$$




Figure 8: Entanglement entropy classification histograms for re-uploading depths $L = 1$ to $L = 4$, where the corresponding test accuracies are 0.49, 0.84, 0.92, and 0.93, respectively. The x -axis shows entanglement entropy and the y -axis the sample count. Blue bars indicate predictions of class 1 and red bars predictions of class 0.




Figure 9: Two classification datasets consisting of unit-norm Bloch vectors (r_1, r_2, r_3) . Red squares and blue circles correspond to classes 0 and 1, respectively.

for computational basis states $|i\rangle$, $|j\rangle$, $|k\rangle$, and $|l\rangle$. We sampled 2-qubit pure states according to the Haar measure [51] and chose A to be the first qubit. Then we assigned label 1 to states with entanglement entropy ≥ 0.3 , which yields approximately balanced classes. We generated a training set of 1000 states and a test set of 500 states, and trained the unrestricted re-uploading circuit in Figure 2a to perform the classification. Results are shown in Figures 6b and 8.

5.4 Classification on the Bloch Sphere

In this section, we test the circuit’s ability to learn a different nonlinear decision boundary by classifying pure single-qubit states on the Bloch sphere. We sampled Bloch vectors uniformly from the unit sphere and assigned label 1 to states satisfying either “ $|r_3| \geq 0.5$ ” (see Figure 9a) or “ $r_3 \geq 0.5$ or $-0.5 \leq r_3 < 0$ ” (see Figure 9b). Each labeling rule yields balanced classes. We generated a training set of 1000 states and a test set of 500 states, and trained the unrestricted re-uploading circuit in Figure 2a to perform the classification. As shown in Figure 10, an accuracy of 0.99 is achieved for (i) Figure 9a with $L = 2$ and (ii) Figure 9b with $L = 3$. This is consistent with the fact that (i) two re-uploading layers suffice to implement $f(r_1, r_2, r_3) = r_3^2$ (with threshold 0.25 for prediction) and (ii) three




Figure 10: Bloch sphere classification histograms for (a) Figure 9a with $L = 1$, (b) Figure 9a with $L = 2$, (c) Figure 9b with $L = 2$, and (d) Figure 9b with $L = 3$, where the corresponding test accuracies are 0.47, 0.99, 0.53, and 0.99, respectively. The x -axis shows r_3 and the y -axis the sample count. Blue bars indicate predictions of class 1 and red bars predictions of class 0.

re-uploading layers suffice to implement $f(r_1, r_2, r_3) = r_3(r_3 - 0.5)(r_3 + 0.5)$ (with threshold 0 for prediction).

6 Discussion and outlook

There exists a fundamental trade-off between expressivity and resource overhead in quantum data re-uploading. On one end, performing successive uploads of a single copy of ρ interleaved with an $(n + 1)$ -qubit unitary reduces qubit count. On the other end, uploading L copies in parallel and applying an $(Ln + 1)$ -qubit unitary maximizes expressivity at the cost of hardware scalability. Characterizing the optimal balance between these extremes remains a challenge.

From a trainability standpoint, the re-uploading ansatz may offer a relative practical advantage over parallel multicopy ansatzes: by keeping the register size small, it might exhibit less severe barren-plateau issues known to worsen with a larger number of qubits [52]. Nonetheless, a detailed ansatz-dependent landscape analysis is essential to guide implementation in near-term devices.

While our universality proof focuses on a single-qubit signal register, extending the framework to the multi-qubit case could enhance model capabilities. A systematic study of how expressivity scales with the size of the signal register, and whether this yields more efficient approximations for particular function classes, is a compelling direction for future work. Moreover, our expressivity analysis was confined to a narrow family of ansatz-agnostic operations involving alternating single-qubit rotations and controlled unitaries, together with specific constructions for polynomial realization. A comprehensive investigation of more general and practical operations remains largely open.

Finally, identifying concrete learning tasks where re-uploading outperforms traditional tomography-plus-post-processing pipelines will be key to demonstrating practical advantage. This entails designing optimal ansatzes for re-uploading circuits tailored to specific quantum datasets.

Author contributions

H.C. led the project, conceived the main idea, developed the main proofs, and wrote the initial draft of the manuscript. D.K.P. established the connection to DQC1, and assisted with physical interpretation and experiments. J.L. helped structure the manuscript, reviewed the text, and assisted in identifying and correcting mathematical errors.

Acknowledgments

This work is in part supported by the National Research Foundation of Korea (NRF, RS-2024-00451435 (10%), RS-2024-00413957 (10%), RS-2023-NR119931 (10%), RS-2025-02309510 (10%)), Institute of Information & Communications Technology Planning & Evaluation (IITP, RS-2021-II212068 (10%), RS-2025-02305453 (10%), RS-2025-02273157 (10%), RS-2025-25442149 (10%), RS-2021-II211343 (10%), No. 2019-0-00003 (10%); Research and Development of Core Technologies for Programming, Running, Implementing and Validating of Fault-Tolerant Quantum Computing System) grant funded by the Ministry of Science and ICT (MSIT), Institute of New Media and Communications (INMAC), and the BK21 FOUR program of the Education, Artificial Intelligence Graduate School Program (Seoul National University), and Research Program for Future ICT Pioneers, Seoul National University in 2025.

References

- [1] Maria Schuld, Ilya Sinayskiy, and Francesco Petruccione. “An introduction to quantum machine learning”. *Contemporary Physics* **56**, 172–185 (2015).
- [2] Jacob Biamonte, Peter Wittek, Nicola Pancotti, Patrick Rebentrost, Nathan Wiebe, and Seth Lloyd. “Quantum machine learning”. *Nature* **549**, 195–202 (2017).
- [3] Nikolaj Moll, Panagiotis Barkoutsos, Lev S Bishop, Jerry M Chow, Andrew Cross, Daniel J Egger, Stefan Filipp, Andreas Fuhrer, Jay M Gambetta, Marc Ganzhorn, et al. “Quantum optimization using variational algorithms on near-term quantum devices”. *Quantum Science and Technology* **3**, 030503 (2018).
- [4] Hsin-Yuan Huang, Michael Broughton, Masoud Mohseni, Ryan Babbush, Sergio Boixo, Hartmut Neven, and Jarrod R McClean. “Power of data in quantum machine learning”. *Nature Communications* **12**, 2631 (2021).
- [5] Yunchao Liu, Srinivasan Arunachalam, and Kristan Temme. “A rigorous and robust quantum speed-up in supervised machine learning”. *Nature Physics* **17**, 1013–1017 (2021).
- [6] Aram W Harrow, Avinatan Hassidim, and Seth Lloyd. “Quantum algorithm for linear systems of equations”. *Physical Review Letters* **103**, 150502 (2009).
- [7] Seth Lloyd, Masoud Mohseni, and Patrick Rebentrost. “Quantum algorithms for supervised and unsupervised machine learning”. *arXiv preprint arXiv:1307.0411* (2013).
- [8] Seth Lloyd, Masoud Mohseni, and Patrick Rebentrost. “Quantum principal component analysis”. *Nature Physics* **10**, 631–633 (2014).
- [9] Marcello Benedetti, Erika Lloyd, Stefan Sack, and Mattia Fiorentini. “Parameterized quantum circuits as machine learning models”. *Quantum Science and Technology* **4**, 043001 (2019).
- [10] Maria Schuld, Ville Bergholm, Christian Gogolin, Josh Izaac, and Nathan Killoran. “Evaluating analytic gradients on quantum hardware”. *Physical Review A* **99**, 032331 (2019).
- [11] Maria Schuld and Nathan Killoran. “Quantum machine learning in feature Hilbert spaces”. *Physical Review Letters* **122**, 040504 (2019).

- [12] Sukin Sim, Peter D Johnson, and Alán Aspuru-Guzik. “Expressibility and entangling capability of parameterized quantum circuits for hybrid quantum-classical algorithms”. *Advanced Quantum Technologies* **2**, 1900070 (2019).
- [13] Carsten Blank, Daniel K Park, June-Koo Kevin Rhee, and Francesco Petruccione. “Quantum classifier with tailored quantum kernel”. *npj Quantum Information* **6**, 41 (2020).
- [14] Marco Cerezo, Andrew Arrasmith, Ryan Babbush, Simon C Benjamin, Suguru Endo, Keisuke Fujii, Jarrod R McClean, Kosuke Mitarai, Xiao Yuan, Lukasz Cincio, et al. “Variational quantum algorithms”. *Nature Reviews Physics* **3**, 625–644 (2021).
- [15] Maria Schuld and Francesco Petruccione. “Quantum models as kernel methods”. In *Machine Learning with Quantum Computers*. Pages 217–245. Springer (2021).
- [16] Sofiene Jerbi, Lukas J Fiderer, Hendrik Poulsen Nautrup, Jonas M Kübler, Hans J Briegel, and Vedran Dunjko. “Quantum machine learning beyond kernel methods”. *Nature Communications* **14**, 517 (2023).
- [17] Matthias C Caro, Elies Gil-Fuster, Johannes Jakob Meyer, Jens Eisert, and Ryan Sweke. “Encoding-dependent generalization bounds for parametrized quantum circuits”. *Quantum* **5**, 582 (2021).
- [18] Maria Schuld, Ryan Sweke, and Johannes Jakob Meyer. “Effect of data encoding on the expressive power of variational quantum-machine-learning models”. *Physical Review A* **103**, 032430 (2021).
- [19] Maria Schuld. “Supervised quantum machine learning models are kernel methods”. *arXiv preprint arXiv:2101.11020* (2021).
- [20] Guangxi Li, Ruilin Ye, Xuanqiang Zhao, and Xin Wang. “Concentration of data encoding in parameterized quantum circuits”. *Advances in Neural Information Processing Systems* **35**, 19456–19469 (2022).
- [21] Shunsuke Daimon and Yu-ichiro Matsushita. “Quantum circuit generation for amplitude encoding using a transformer decoder”. *Physical Review Applied* **22**, L041001 (2024).
- [22] Javier Gonzalez-Conde, Thomas W Watts, Pablo Rodriguez-Grasa, and Mikel Sanz. “Efficient quantum amplitude encoding of polynomial functions”. *Quantum* **8**, 1297 (2024).
- [23] Tak Hur, Israel F Araujo, and Daniel K Park. “Neural quantum embedding: Pushing the limits of quantum supervised learning”. *Physical Review A* **110**, 022411 (2024).
- [24] Ben Jaderberg, Antonio A Gentile, Youssef Achari Berrada, Elvira Shishenina, and Vincent E Elfving. “Let quantum neural networks choose their own frequencies”. *Physical Review A* **109**, 042421 (2024).
- [25] Minati Rath and Hema Date. “Quantum data encoding: A comparative analysis of classical-to-quantum mapping techniques and their impact on machine learning accuracy”. *EPJ Quantum Technology* **11**, 72 (2024).
- [26] Hongfeng Liu, Tak Hur, Shitao Zhang, Liangyu Che, Xinyue Long, Xiangyu Wang, Keyi Huang, Yu-ang Fan, Yuxuan Zheng, Yufang Feng, et al. “Neural quantum embedding via deterministic quantum computation with one qubit”. *Physical Review Letters* **135**, 080603 (2025).
- [27] Adrián Pérez-Salinas, Alba Cervera-Lierta, Elies Gil-Fuster, and José I Latorre. “Data re-uploading for a universal quantum classifier”. *Quantum* **4**, 226 (2020).

- [28] Adrián Pérez-Salinas, David López-Núñez, Artur García-Sáez, Pol Forn-Díaz, and José I Latorre. “One qubit as a universal approximant”. *Physical Review A* **104**, 012405 (2021).
- [29] Yujin Kim and Daniel K. Park. “Expressivity of deterministic quantum computation with one qubit”. *Physical Review A* **111**, 022429 (2025).
- [30] Guang Hao Low and Isaac L Chuang. “Optimal Hamiltonian simulation by quantum signal processing”. *Physical Review Letters* **118**, 010501 (2017).
- [31] John M Martyn, Zane M Rossi, Andrew K Tan, and Isaac L Chuang. “Grand unification of quantum algorithms”. *PRX Quantum* **2**, 040203 (2021).
- [32] Zane M Rossi and Isaac L Chuang. “Multivariable quantum signal processing (MQSP): prophecies of the two-headed oracle”. *Quantum* **6**, 811 (2022).
- [33] Danial Motlagh and Nathan Wiebe. “Generalized quantum signal processing”. *PRX Quantum* **5**, 020368 (2024).
- [34] Jonas Kübler, Simon Buchholz, and Bernhard Schölkopf. “The inductive bias of quantum kernels”. *Advances in Neural Information Processing Systems* **34**, 12661–12673 (2021).
- [35] Marco Cerezo, Guillaume Verdon, Hsin-Yuan Huang, Lukasz Cincio, and Patrick J Coles. “Challenges and opportunities in quantum machine learning”. *Nature Computational Science* **2**, 567–576 (2022).
- [36] Hsin-Yuan Huang, Michael Broughton, Jordan Cotler, Sitan Chen, Jerry Li, Masoud Mohseni, Hartmut Neven, Ryan Babbush, Richard Kueng, John Preskill, et al. “Quantum advantage in learning from experiments”. *Science* **376**, 1182–1186 (2022).
- [37] Artur K Ekert, Carolina Moura Alves, Daniel KL Oi, Michał Horodecki, Paweł Horodecki, and Leong Chuan Kwek. “Direct estimations of linear and nonlinear functionals of a quantum state”. *Physical Review Letters* **88**, 217901 (2002).
- [38] Todd A. Brun. “Measuring polynomial functions of states”. *Quantum Information & Computation* **4**, 401–408 (2004).
- [39] Hongshun Yao, Yingjian Liu, Tengxiang Lin, and Xin Wang. “Sample-efficient estimation of nonlinear quantum state functions”. *Communications Physics* (2026).
- [40] Ludmila Botelho, Adam Glos, Akash Kundu, Jarosław Adam Mischczak, Özlem Salehi, and Zoltán Zimborás. “Error mitigation for variational quantum algorithms through mid-circuit measurements”. *Physical Review A* **105**, 022441 (2022).
- [41] Sebastian Brandhofer, Ilia Polian, and Kevin Krsulich. “Optimal qubit reuse for near-term quantum computers”. In *2023 IEEE International Conference on Quantum Computing and Engineering (QCE)*. Volume 1, pages 859–869. IEEE (2023).
- [42] Matthew DeCross, Eli Chertkov, Megan Kohagen, and Michael Foss-Feig. “Qubit-reuse compilation with mid-circuit measurement and reset”. *Physical Review X* **13**, 041057 (2023).
- [43] Daiwei Zhu, Gregory D Kahanamoku-Meyer, Laura Lewis, Crystal Noel, Or Katz, Bahaa Harraz, Qingfeng Wang, Andrew Risinger, Lei Feng, Debopriyo Biswas, et al. “Interactive cryptographic proofs of quantumness using mid-circuit measurements”. *Nature Physics* **19**, 1725–1731 (2023).
- [44] Emanuel Knill and Raymond Laflamme. “Power of one bit of quantum information”. *Physical Review Letters* **81**, 5672 (1998).

- [45] Dan Shepherd. “Computation with unitaries and one pure qubit”. [arXiv preprint quant-ph/0608132](#) (2006).
- [46] Lars Hörmander. “The analysis of linear partial differential operators I: Distribution theory and Fourier analysis”. [Springer](#). (2015).
- [47] Adrián Pérez-Salinas, Mahtab Yaghubi Rad, Alice Barthe, and Vedran Dunjko. “Universal approximation of continuous functions with minimal quantum circuits”. [Physical Review Research](#) **7**, 043282 (2025).
- [48] Kosuke Mitarai, Makoto Negoro, Masahiro Kitagawa, and Keisuke Fujii. “Quantum circuit learning”. [Physical Review A](#) **98**, 032309 (2018).
- [49] Hyunho Cha. “Code repository”. <https://github.com/HyunHoCha/reupload> (2026).
- [50] Michael A Nielsen and Isaac L Chuang. “Quantum Computation and Quantum Information”. [Cambridge University Press](#). (2010).
- [51] Antonio Anna Mele. “Introduction to Haar measure tools in quantum information: A beginner’s tutorial”. [Quantum](#) **8**, 1340 (2024).
- [52] Jarrod R McClean, Sergio Boixo, Vadim N Smelyanskiy, Ryan Babbush, and Hartmut Neven. “Barren plateaus in quantum neural network training landscapes”. [Nature Communications](#) **9**, 4812 (2018).
- [53] Adriano Barenco, André Berthiaume, David Deutsch, Artur Ekert, Richard Jozsa, and Chiara Macchiavello. “Stabilization of quantum computations by symmetrization”. [SIAM Journal on Computing](#) **26**, 1541–1557 (1997).
- [54] Harry Buhrman, Richard Cleve, John Watrous, and Ronald De Wolf. “Quantum fingerprinting”. [Physical Review Letters](#) **87**, 167902 (2001).
- [55] Robert R Tucci. “An introduction to Cartan’s KAK decomposition for QC programmers”. [arXiv preprint quant-ph/0507171](#) (2005).
- [56] David Wierichs, Maxwell West, Roy T Forestano, M Cerezo, and Nathan Kilorian. “Recursive Cartan decompositions for unitary synthesis”. [arXiv preprint arXiv:2503.19014](#) (2025).

A Derivation of Eq. (11)

Note that

$$\begin{aligned} \sum_{\beta=1}^{d/2} \langle \psi_{\beta}^{(\alpha)} | W_{\alpha'} | \psi_{\beta}^{(\alpha)} \rangle &= \text{tr} \left(\sum_{\beta=1}^{d/2} | \psi_{\beta}^{(\alpha)} \rangle \langle \psi_{\beta}^{(\alpha)} | W_{\alpha'} \right) = \text{tr}(W_{\alpha}^{+} W_{\alpha'}), \\ \sum_{\gamma=1}^{d/2} \langle \phi_{\gamma}^{(\alpha)} | W_{\alpha'} | \phi_{\gamma}^{(\alpha)} \rangle &= \text{tr} \left(\sum_{\gamma=1}^{d/2} | \phi_{\gamma}^{(\alpha)} \rangle \langle \phi_{\gamma}^{(\alpha)} | W_{\alpha'} \right) = \text{tr}(W_{\alpha}^{-} W_{\alpha'}), \end{aligned}$$

and

$$\begin{aligned} \text{tr}(W_{\alpha}^{+} W_{\alpha'}) + \text{tr}(W_{\alpha}^{-} W_{\alpha'}) &= \text{tr}(W_{\alpha'}) = 0, \\ \text{tr}(W_{\alpha}^{+} W_{\alpha'}) - \text{tr}(W_{\alpha}^{-} W_{\alpha'}) &= \text{tr}(W_{\alpha} W_{\alpha'}) = d \cdot \delta_{\alpha, \alpha'}. \end{aligned}$$

Then we have

$$\begin{aligned}
\tau^{(l)} &= \sum_{\beta=1}^{d/2} K_{\beta}^{(\alpha+)} \left(\tilde{\tau}^{(l)} \otimes \rho \right) K_{\beta}^{(\alpha+)\dagger} + \sum_{\gamma=1}^{d/2} K_{\gamma}^{(\alpha-)} \left(\tilde{\tau}^{(l)} \otimes \rho \right) K_{\gamma}^{(\alpha-)\dagger} \\
&= \frac{1}{2d} \left(\mathbb{I}_2 + \tilde{r}_1^{(l)} X + \tilde{r}_2^{(l)} Y + \tilde{r}_3^{(l)} Z \right) \sum_{\beta=1}^{d/2} \left(1 + \sum_{\alpha'=1}^{d^2-1} \lambda_{\alpha'} \left\langle \psi_{\beta}^{(\alpha)} \middle| W_{\alpha'} \middle| \psi_{\beta}^{(\alpha)} \right\rangle \right) \\
&\quad + \frac{1}{2d} \left(\mathbb{I}_2 + \tilde{r}_1^{(l)} X - \tilde{r}_2^{(l)} Y - \tilde{r}_3^{(l)} Z \right) \sum_{\gamma=1}^{d/2} \left(1 + \sum_{\alpha'=1}^{d^2-1} \lambda_{\alpha'} \left\langle \phi_{\gamma}^{(\alpha)} \middle| W_{\alpha'} \middle| \phi_{\gamma}^{(\alpha)} \right\rangle \right) \\
&= \frac{1}{2d} \left(\mathbb{I}_2 + \tilde{r}_1^{(l)} X \right) \left(d + \sum_{\alpha'=1}^{d^2-1} \lambda_{\alpha'} \left[\text{tr}(W_{\alpha}^+ W_{\alpha'}) + \text{tr}(W_{\alpha}^- W_{\alpha'}) \right] \right) \\
&\quad + \frac{1}{2d} \left(\tilde{r}_2^{(l)} Y + \tilde{r}_3^{(l)} Z \right) \sum_{\alpha'=1}^{d^2-1} \lambda_{\alpha'} \left[\text{tr}(W_{\alpha}^+ W_{\alpha'}) - \text{tr}(W_{\alpha}^- W_{\alpha'}) \right] \\
&= \frac{1}{2} \left(\mathbb{I}_2 + \tilde{r}_1^{(l)} X \right) + \frac{1}{2d} \left(\tilde{r}_2^{(l)} Y + \tilde{r}_3^{(l)} Z \right) \sum_{\alpha'=1}^{d^2-1} \lambda_{\alpha'} d \cdot \delta_{\alpha, \alpha'} \\
&= \frac{1}{2} \left(\mathbb{I}_2 + \tilde{r}_1^{(l)} X \right) + \frac{1}{2} \left(\tilde{r}_2^{(l)} Y + \tilde{r}_3^{(l)} Z \right) \lambda_{\alpha} \\
&= \frac{1}{2} \left(\mathbb{I}_2 + \tilde{r}_1^{(l)} X + [\lambda_{\alpha} \tilde{r}_2^{(l)}] Y + [\lambda_{\alpha} \tilde{r}_3^{(l)}] Z \right).
\end{aligned}$$

B Derivation of Eq. (14)

$$\begin{aligned}
U_{\alpha, i}(\delta) &= \sum_{q=0}^{\infty} \frac{(-i\delta)^q}{q!} H_{\alpha, i}^q \\
&= \sum_{m=0}^{\infty} \frac{(-i\delta)^{2m}}{(2m)!} H_{\alpha, i}^{2m} + \sum_{m=0}^{\infty} \frac{(-i\delta)^{2m+1}}{(2m+1)!} H_{\alpha, i}^{2m+1} \\
&= \left(\sum_{m=0}^{\infty} \frac{(-1)^m \delta^{2m}}{(2m)!} \right) (\mathbb{I}_2)_A \otimes (\mathbb{I}_d)_B - i \left(\sum_{m=0}^{\infty} \frac{(-1)^m \delta^{2m+1}}{(2m+1)!} \right) \sigma_A^{(i)} \otimes (W_{\alpha})_B \\
&= \cos \delta (\mathbb{I}_2)_A \otimes (\mathbb{I}_d)_B - i \sin \delta \sigma_A^{(i)} \otimes (W_{\alpha})_B
\end{aligned}$$

C Proof of Theorem 2

Proof. First, the generalized Bloch coordinate map is faithful. Indeed, $\rho \mapsto \boldsymbol{\lambda}(\rho)$ is an affine homeomorphism from $\mathcal{D}(\mathcal{H}_d)$ onto the compact physical Bloch body Λ_d . Hence the target induces a continuous function $f_{\Lambda} : \Lambda_d \rightarrow \mathbb{R}$, defined by $f_{\Lambda}(\boldsymbol{\lambda}(\rho)) = F(\rho)$. By the Tietze extension theorem, extend f_{Λ} to a continuous function $f : [-1, 1]^m \rightarrow \mathbb{R}$. Choose constants $a > 0$ and $b \in \mathbb{R}$ such that

$$g(\boldsymbol{\lambda}) = \frac{f(\boldsymbol{\lambda}) - b}{a} \in [-1, 1] \quad \text{for all } \boldsymbol{\lambda} \in [-1, 1]^m.$$

In particular, one may choose a to be the right-hand side of Eq. (25). After the affine change of variables $\mathbf{x} = (\boldsymbol{\lambda} + \mathbf{1})/2 \in [0, 1]^m$, define $g^{\sharp}(\mathbf{x}) = g(2\mathbf{x} - \mathbf{1})$.

We now invoke the classical re-uploading universality. The underlying single-qubit classical data re-uploading universality in $\|\cdot\|_\infty$ is established in Ref. [47]. It implies that, for any $\eta_{\text{cl}} > 0$, there exist a finite single-qubit classical re-uploading circuit $\mathcal{C}_\theta$ and an observable $O_{\text{cl}} \in \text{Herm}(\mathcal{H}_2)$ with $\|O_{\text{cl}}\|_\infty \leq 1$ such that

$$\sup_{\mathbf{x} \in [0,1]^m} \left| \text{tr}(O_{\text{cl}} \mathcal{C}_\theta(2\mathbf{x} - \mathbf{1})(|0\rangle\langle 0|)) - g^\sharp(\mathbf{x}) \right| < \eta_{\text{cl}}. \quad (28)$$

The circuit uses single-qubit rotations whose angles are affine functions of the coordinates of $\mathbf{x}$. Since $x_\alpha = (\lambda_\alpha + 1)/2$, each factor depending on x_α can be written as an input-independent single-qubit rotation composed with a data-dependent rotation of the form $\exp(-it\lambda_\alpha\sigma^{(i)})$, after possibly changing the real parameter t and absorbing global phases. Therefore the circuit can be written in the form of Eq. (24).

It remains to show that each data-dependent classical rotation in $\mathcal{C}_\theta$ is implementable by quantum uploads of ρ . Fix α, i , and δ with $|\delta| \leq 1$. The weak collision channel derived in Eq. (17) is

$$\Phi_{\alpha,i}^{(\delta)}(\tau) = \cos^2 \delta \tau + \sin^2 \delta \sigma^{(i)} \tau \sigma^{(i)} - i\lambda_\alpha(\rho) \sin \delta \cos \delta [\sigma^{(i)}, \tau],$$

whereas the ideal classical rotation channel in Eq. (21) is

$$\mathcal{R}_{\alpha,i}^{(\delta)}(\tau) = \cos^2(\delta\lambda_\alpha)\tau + \sin^2(\delta\lambda_\alpha)\sigma^{(i)}\tau\sigma^{(i)} - i\sin(\delta\lambda_\alpha)\cos(\delta\lambda_\alpha)[\sigma^{(i)}, \tau].$$

Using $|\lambda_\alpha| \leq 1$, $\|\tau\|_1 = 1$, $\|\sigma^{(i)}\tau\sigma^{(i)}\|_1 = 1$, and $\|[\sigma^{(i)}, \tau]\|_1 \leq 2$, we obtain the uniform single-step bound

$$\sup_{\rho \in \mathcal{D}(\mathcal{H}_d)} \sup_{\tau \in \mathcal{D}(\mathcal{H}_2)} \left\| \Phi_{\alpha,i}^{(\delta)}(\tau) - \mathcal{R}_{\alpha,i}^{(\delta)}(\tau) \right\|_1 \leq 6\delta^2. \quad (29)$$

Now choose an integer $N \geq |t|$ and set $\delta = t/N$. Since trace norm is contractive under completely positive trace-preserving maps, and since the ideal rotations satisfy $(\mathcal{R}_{\alpha,i}^{(t/N)})^N = \mathcal{R}_{\alpha,i}^{(t)}$, a standard telescoping estimate gives

$$\sup_{\rho, \tau} \left\| (\Phi_{\alpha,i}^{(t/N)})^N(\tau) - \mathcal{R}_{\alpha,i}^{(t)}(\tau) \right\|_1 \leq N \cdot 6(t/N)^2 = \frac{6t^2}{N}.$$

Applying the same argument over the M data-dependent rotations in $\mathcal{C}_\theta$, while noting that the input-independent channels $\mathcal{V}^{(j)}$ are isometries with respect to the trace norm, gives Eq. (26).

Finally choose $\eta_{\text{cl}} = \varepsilon/(2a)$ in Eq. (28) and choose the integers $N_1, \dots, N_M$ such that

$$6 \sum_{j=1}^M \frac{t_j^2}{N_j} < \frac{\varepsilon}{2a}.$$

Let $\tau_A^{\text{out}}(\rho) = \hat{\mathcal{C}}_{\theta, \mathbf{N}}^{\rho}(|0\rangle\langle 0|)$, set $O_{\text{out}} = aO_{\text{cl}}$, and set $b_{\text{out}} = b$. Then, for all $\rho \in \mathcal{D}(\mathcal{H}_d)$,

$$\begin{aligned} \left| b_{\text{out}} + \text{tr}(O_{\text{out}} \tau_A^{\text{out}}(\rho)) - F(\rho) \right| &\leq a \left| \text{tr}(O_{\text{cl}} \tau_A^{\text{out}}(\rho)) - \text{tr}(O_{\text{cl}} \mathcal{C}_\theta(\boldsymbol{\lambda}(\rho))(|0\rangle\langle 0|)) \right| \\ &\quad + a \left| \text{tr}(O_{\text{cl}} \mathcal{C}_\theta(\boldsymbol{\lambda}(\rho))(|0\rangle\langle 0|)) - g(\boldsymbol{\lambda}(\rho)) \right| \\ &< a \cdot \frac{\varepsilon}{2a} + a \cdot \frac{\varepsilon}{2a} = \varepsilon. \end{aligned}$$

Taking the supremum over ρ proves the claimed uniform approximation guarantee. The readout is an affine combination of single-qubit Pauli operators because every Hermitian $O_{\text{out}} \in \text{Herm}(\mathcal{H}_2)$ is a real linear combination of $\mathbb{I}_2, X, Y, Z$, and the identity contribution is absorbed into b_{out} . $\square$

D Details of the purity test




Figure 11: The swap test circuit [53, 54], typically used to estimate inner products between pure states, also serves as a tool for estimating the purity of a quantum state ρ .

Figure 11 shows that applying the swap test to two copies of ρ enables evaluation of the purity of ρ . In more detail, write

$$\rho = \sum_i \mu_i |\psi_i\rangle \langle \psi_i|,$$

where $\mu_i \geq 0$ and $\{|\psi_i\rangle\}_i$ is an orthonormal basis. Then

$$\rho^{\otimes 2} = \sum_{i,j} \mu_i \mu_j |\psi_i\rangle \langle \psi_i| \otimes |\psi_j\rangle \langle \psi_j|$$

and

$$\begin{aligned} P(\text{First qubit} = 0 \mid \rho^{\otimes 2}) &= \sum_{i,j} \mu_i \mu_j P(\text{First qubit} = 0 \mid |\psi_i\rangle \otimes |\psi_j\rangle) \\ &= \sum_{i,j} \mu_i \mu_j \frac{1}{2} (1 + \delta_{i,j}) \\ &= \frac{1}{2} \sum_i \mu_i^2 + \frac{1}{2} \left(\sum_i \mu_i \right)^2 \\ &= \frac{1}{2} + \frac{1}{2} \text{tr}(\rho^2). \end{aligned}$$

The expectation value $\langle Z \rangle$ on system A is equal to the purity $\text{tr}(\rho^2)$.

E Proof of Observation 1

Proof. We show that no collection (U, V, O) exists such that $U, V \in \text{SU}(4)$, $O \in \text{Herm}(\mathcal{H}_2)$, and $\langle O_A \rangle_{\sigma_A} = \text{tr}(\rho^2)$, where

$$\sigma_A = \text{tr}_{B_2} \left(V_{AB_2} \left(\text{tr}_{B_1} \left(U_{AB_1} (|0\rangle \langle 0|_A \otimes \rho_{B_1}) U_{AB_1}^\dagger \right) \otimes \rho_{B_2} \right) V_{AB_2}^\dagger \right).$$

The proof compares two facts about the effective two-copy observable: the re-uploading form forces a singular Pauli correlation matrix, whereas purity forces a nonsingular one. Write

$$K_k = ((\mathbb{I}_2)_A \otimes \langle k_{B_1} |) U_{AB_1} (|0_A\rangle \otimes (\mathbb{I}_2)_{B_1}) \quad (k = 0, 1).$$

After the first upload, the state of system A is given by

$$\sum_{k=0}^1 K_k \rho_{B_1} K_k^\dagger.$$

For the second upload set

$$J_k = V_{AB_2}(K_k \otimes (\mathbb{I}_2)_{B_2}), \quad \tilde{O}_{AB_2} = O_A \otimes (\mathbb{I}_2)_{B_2}.$$

Then

$$\begin{aligned} \langle O_A \rangle_{\sigma_A} &= \text{tr}_{AB_2} \left(\left[\sum_{k=0}^1 J_k (\rho_{B_1} \otimes \rho_{B_2}) J_k^\dagger \right] \tilde{O}_{AB_2} \right) \\ &= \text{tr}_{B_1 B_2} \left((\rho_{B_1} \otimes \rho_{B_2}) \left[\sum_{k=0}^1 J_k^\dagger \tilde{O}_{AB_2} J_k \right] \right). \end{aligned}$$

For $M \in \text{Herm}(\mathcal{H}_2 \otimes \mathcal{H}_2)$, define $\text{corr}(M) \in \mathbb{R}^{3 \times 3}$ such that

$$\text{corr}(M)_{ij} = \frac{1}{2} \text{tr} \left(M \left(\sigma^{(i)} \otimes \sigma^{(j)} \right) \right).$$

In what follows, we show that

$$\det \left(\text{corr} \left(\sum_{k=0}^1 J_k^\dagger \tilde{O}_{AB_2} J_k \right) \right) = 0.$$

For notational simplicity, we omit system subscripts and show that $\det(T) = 0$, where

$$\begin{aligned} K_k &= (\mathbb{I}_2 \otimes \langle k |) U(|0\rangle \otimes \mathbb{I}_2), \quad J_k = V(K_k \otimes \mathbb{I}_2), \quad \tilde{O} = O \otimes \mathbb{I}_2, \\ \text{and } T &= \text{corr} \left(\sum_{k=0}^1 J_k^\dagger \tilde{O} J_k \right). \end{aligned}$$

Let

$$Q = V^\dagger \tilde{O} V$$

and define

$$\Lambda(\cdot) = \sum_{k=0}^1 K_k(\cdot) K_k^\dagger.$$

Then

$$\begin{aligned} T_{ij} &= \frac{1}{2} \text{tr} \left(\sum_{k=0}^1 (K_k^\dagger \otimes \mathbb{I}_2) Q (K_k \otimes \mathbb{I}_2) (\sigma^{(i)} \otimes \sigma^{(j)}) \right) \\ &= \frac{1}{2} \text{tr} \left(Q \sum_{k=0}^1 (K_k \otimes \mathbb{I}_2) (\sigma^{(i)} \otimes \sigma^{(j)}) (K_k^\dagger \otimes \mathbb{I}_2) \right) \\ &= \frac{1}{2} \text{tr} \left(Q \left(\sum_{k=0}^1 K_k \sigma^{(i)} K_k^\dagger \otimes \sigma^{(j)} \right) \right) \\ &= \frac{1}{2} \text{tr} \left(Q \left(\Lambda(\sigma^{(i)}) \otimes \sigma^{(j)} \right) \right). \end{aligned}$$

But note that $\Lambda(\sigma^{(i)})$ is traceless for $i = 1, 2, 3$ because

$$\text{tr} \left(\Lambda(\sigma^{(i)}) \right) = \text{tr} \left(\sigma^{(i)} \sum_{k=0}^1 K_k^\dagger K_k \right) = \text{tr}(\sigma^{(i)}) = 0.$$

Let $C = \text{corr}(Q)$ and define $\tilde{C} \in \mathbb{R}^{3 \times 3}$ such that

$$\tilde{C}_{\mu i} = \frac{1}{2} \text{tr} \left(\sigma^{(\mu)} \Lambda \left(\sigma^{(i)} \right) \right), \quad \mu, i \in \{1, 2, 3\}.$$

Then $\Lambda \left(\sigma^{(i)} \right)$ for $i = 1, 2, 3$ can be expressed as

$$\begin{aligned} \Lambda \left(\sigma^{(i)} \right) &= \sum_{\mu=1}^3 \frac{1}{2} \text{tr} \left(\sigma^{(\mu)} \Lambda \left(\sigma^{(i)} \right) \right) \sigma^{(\mu)} \\ &= \sum_{\mu=1}^3 \tilde{C}_{\mu i} \sigma^{(\mu)}, \end{aligned}$$

hence

$$\begin{aligned} T_{ij} &= \frac{1}{2} \text{tr} \left(Q \left(\sum_{\mu=1}^3 \tilde{C}_{\mu i} \sigma^{(\mu)} \otimes \sigma^{(j)} \right) \right) \\ &= \sum_{\mu=1}^3 \tilde{C}_{\mu i} \frac{1}{2} \text{tr} \left(Q \left(\sigma^{(\mu)} \otimes \sigma^{(j)} \right) \right) \\ &= \sum_{\mu=1}^3 \tilde{C}_{\mu i} C_{\mu j} \\ &= (\tilde{C}^T C)_{ij}. \end{aligned}$$

Therefore, it suffices to show that $\det(C) = \det(\text{corr}(V^\dagger \tilde{O} V)) = 0$. Note that any $V \in \text{SU}(4)$ allows the *KAK decomposition*

$$V = (\mathfrak{A}_1 \otimes \mathfrak{B}_1) e^{i\mathbf{k} \cdot \Sigma} (\mathfrak{A}_2 \otimes \mathfrak{B}_2),$$

where $\mathfrak{A}_1, \mathfrak{A}_2, \mathfrak{B}_1, \mathfrak{B}_2 \in \text{SU}(2)$, $\mathbf{k} = (k_1, k_2, k_3) \in \mathbb{R}^3$, and

$$\Sigma = \left(\sigma^{(1)} \otimes \sigma^{(1)}, \sigma^{(2)} \otimes \sigma^{(2)}, \sigma^{(3)} \otimes \sigma^{(3)} \right)$$

[55, 56]. Also, every $\mathfrak{U} \in \text{SU}(2)$ is associated with $\mathfrak{R}(\mathfrak{U}) \in \text{SO}(3)$ such that

$$\mathfrak{U} \sigma^{(i)} \mathfrak{U}^\dagger = \sum_{j=1}^3 \mathfrak{R}(\mathfrak{U})_{ij} \sigma^{(j)}, \quad i = 1, 2, 3.$$

Then

$$\begin{aligned} &\text{tr} \left(V^\dagger \tilde{O} V \left(\sigma^{(i)} \otimes \sigma^{(j)} \right) \right) \\ &= \text{tr} \left(\left(\mathfrak{A}_2^\dagger \otimes \mathfrak{B}_2^\dagger \right) e^{-i\mathbf{k} \cdot \Sigma} \left(\mathfrak{A}_1^\dagger \otimes \mathfrak{B}_1^\dagger \right) (O \otimes \mathbb{I}_2) (\mathfrak{A}_1 \otimes \mathfrak{B}_1) e^{i\mathbf{k} \cdot \Sigma} (\mathfrak{A}_2 \otimes \mathfrak{B}_2) \left(\sigma^{(i)} \otimes \sigma^{(j)} \right) \right) \\ &= \text{tr} \left(e^{-i\mathbf{k} \cdot \Sigma} (O' \otimes \mathbb{I}_2) e^{i\mathbf{k} \cdot \Sigma} \left(\mathfrak{A}_2 \sigma^{(i)} \mathfrak{A}_2^\dagger \otimes \mathfrak{B}_2 \sigma^{(j)} \mathfrak{B}_2^\dagger \right) \right) \\ &= \sum_{p,q=1}^3 \mathfrak{R}(\mathfrak{A}_2)_{ip} \mathfrak{R}(\mathfrak{B}_2)_{jq} \text{tr} \left(e^{-i\mathbf{k} \cdot \Sigma} (O' \otimes \mathbb{I}_2) e^{i\mathbf{k} \cdot \Sigma} \left(\sigma^{(p)} \otimes \sigma^{(q)} \right) \right) \\ &= 2 \left(\mathfrak{R}(\mathfrak{A}_2) \text{corr} \left(e^{-i\mathbf{k} \cdot \Sigma} (O' \otimes \mathbb{I}_2) e^{i\mathbf{k} \cdot \Sigma} \right) \mathfrak{R}(\mathfrak{B}_2)^T \right)_{ij}, \end{aligned}$$

where $O' = \mathfrak{A}_1^\dagger O \mathfrak{A}_1$. The local factors only left- and right-multiply by $\text{SO}(3)$ matrices, so they do not affect whether the determinant vanishes. Now it suffices to show that

$$\det\left(\text{corr}\left(e^{-i\mathbf{k}\cdot\boldsymbol{\Sigma}}(O' \otimes \mathbb{I}_2)e^{i\mathbf{k}\cdot\boldsymbol{\Sigma}}\right)\right) = 0.$$

Since $[\sigma^{(i)} \otimes \sigma^{(i)}, \sigma^{(j)} \otimes \sigma^{(j)}] = 0$ for any (i, j) and each $\sigma^{(i)} \otimes \sigma^{(i)}$ squares to the identity, $e^{i\mathbf{k}\cdot\boldsymbol{\Sigma}}$ can be factorized as

$$e^{i\mathbf{k}\cdot\boldsymbol{\Sigma}} = \left[\cos k_1 \mathbb{I}_4 + i \sin k_1 \left(\sigma^{(1)} \otimes \sigma^{(1)} \right) \right] \left[\cos k_2 \mathbb{I}_4 + i \sin k_2 \left(\sigma^{(2)} \otimes \sigma^{(2)} \right) \right] \left[\cos k_3 \mathbb{I}_4 + i \sin k_3 \left(\sigma^{(3)} \otimes \sigma^{(3)} \right) \right].$$

A direct calculation yields

$$\begin{aligned} \text{corr}\left(e^{-i\mathbf{k}\cdot\boldsymbol{\Sigma}} \left(\sigma^{(1)} \otimes \mathbb{I}_2 \right) e^{i\mathbf{k}\cdot\boldsymbol{\Sigma}}\right) &= 2 \left[\cos(2k_2) \sin(2k_3) \mathbf{e}_2 \mathbf{e}_3^T - \sin(2k_2) \cos(2k_3) \mathbf{e}_3 \mathbf{e}_2^T \right], \\ \text{corr}\left(e^{-i\mathbf{k}\cdot\boldsymbol{\Sigma}} \left(\sigma^{(2)} \otimes \mathbb{I}_2 \right) e^{i\mathbf{k}\cdot\boldsymbol{\Sigma}}\right) &= -2 \left[\cos(2k_1) \sin(2k_3) \mathbf{e}_1 \mathbf{e}_3^T - \sin(2k_1) \cos(2k_3) \mathbf{e}_3 \mathbf{e}_1^T \right], \\ \text{corr}\left(e^{-i\mathbf{k}\cdot\boldsymbol{\Sigma}} \left(\sigma^{(3)} \otimes \mathbb{I}_2 \right) e^{i\mathbf{k}\cdot\boldsymbol{\Sigma}}\right) &= 2 \left[\cos(2k_1) \sin(2k_2) \mathbf{e}_1 \mathbf{e}_2^T - \sin(2k_1) \cos(2k_2) \mathbf{e}_2 \mathbf{e}_1^T \right]. \end{aligned}$$

One can verify that any linear combination of these three matrices has determinant zero. Since O' can be expressed as a linear combination of $\{\sigma^{(i)}\}_{i=0}^3$ and $\text{corr}\left(e^{-i\mathbf{k}\cdot\boldsymbol{\Sigma}} \left(\sigma^{(0)} \otimes \mathbb{I}_2 \right) e^{i\mathbf{k}\cdot\boldsymbol{\Sigma}}\right)$ evaluates to the zero matrix (hence contributing nothing to $\text{corr}(\cdot)$), we have

$$\det\left(\text{corr}\left(e^{-i\mathbf{k}\cdot\boldsymbol{\Sigma}}(O' \otimes \mathbb{I}_2)e^{i\mathbf{k}\cdot\boldsymbol{\Sigma}}\right)\right) = 0.$$

Meanwhile, an element-wise expansion of $\rho \otimes \rho$ reveals that if $\hat{O} \in \text{Herm}(\mathcal{H}_2 \otimes \mathcal{H}_2)$ satisfies

$$\text{tr}((\rho \otimes \rho) \hat{O}) = \text{tr}(\rho^2)$$

for all $\rho \in \mathcal{D}(\mathcal{H}_2)$, then it must take the form

$$\begin{aligned} \hat{O} &= \text{SWAP} + c_1 i \left(\mathbf{e}_2 \mathbf{e}_3^T - \mathbf{e}_3 \mathbf{e}_2^T \right) + c_2 \left(\mathbf{e}_2 \mathbf{e}_2^T - \mathbf{e}_3 \mathbf{e}_3^T \right) \\ &\quad + (c_3 + c_4 i) \left(\mathbf{e}_1 \mathbf{e}_2^T - \mathbf{e}_1 \mathbf{e}_3^T \right) + (c_3 - c_4 i) \left(\mathbf{e}_2 \mathbf{e}_1^T - \mathbf{e}_3 \mathbf{e}_1^T \right) \\ &\quad + (c_5 + c_6 i) \left(\mathbf{e}_2 \mathbf{e}_4^T - \mathbf{e}_3 \mathbf{e}_4^T \right) + (c_5 - c_6 i) \left(\mathbf{e}_4 \mathbf{e}_2^T - \mathbf{e}_4 \mathbf{e}_3^T \right) \end{aligned}$$

for some $c_1, \dots, c_6 \in \mathbb{R}$. This gives

$$\text{corr}(\hat{O}) = \begin{pmatrix} 1 & c_1 & -(c_3 + c_5) \\ -c_1 & 1 & c_4 + c_6 \\ c_3 + c_5 & -(c_4 + c_6) & 1 \end{pmatrix},$$

which has determinant $1 + c_1^2 + (c_3 + c_5)^2 + (c_4 + c_6)^2 > 0$. Since the same effective observable cannot have both $\det = 0$ and $\det > 0$, this concludes the proof of Observation 1. $\square$