

**Chavis Srichan**

Department of Computer Engineering,
Khon Kaen University, Khon Kaen, Thailand
ORCID: 0000-0002-4772-3543
chavis@kku.ac.th

Isamu Ohnishi

Graduate School of Integrated Sciences for Life,
Hiroshima University, Higashi-Hiroshima, Japan
ORCID: 0009-0009-5716-6216
isamu_o@toki.waseda.jp (corresponding author)

EXACT ALGEBRAIC REDUCTION OF QUANTUM NEURAL CIRCUITS: OBSERVABLE-DEPENDENT RENORMALIZATION AND CURVATURE-BASED RELIABILITY CONTROL

Abstract. The computational cost and operational reliability of parameterized quantum circuits depend on which outputs an information system must preserve. This paper develops an observable-centred workflow linking quantum neural circuits, Lie and Krylov closure, exact algebraic reduction, and a proposed geometric reliability monitor. The dynamical model is a finite-dimensional Markovian quantum system with a declared set of target observables. We distinguish three objects that must not be conflated: the minimal linear observable space, a closed Pauli-coordinate space used for classical simulation, and an associative operator algebra supporting a completely positive trace-preserving reduced quantum model. A direct invariance argument establishes exact preservation of the declared outputs; an explicit sufficient construction specifies when this linear result also admits a physical quantum realization. A structured XX chain with local dephasing illustrates the reduction from an ambient operator space of dimension 4^n to a Pauli-coordinate space of dimension $n(2n-1)$, including 45 rather than 1024 coordinates for five qubits. A separately specified numerical check compares reduced and full evolution, while a generic interacting model demonstrates that the same coordinate-closure procedure need not compress. Quantum Fisher and Berry geometry are introduced with additional sufficiency and smoothness conditions rather than assumed to survive arbitrary output compression. An operational decomposition separates gauge, locally output-insensitive, and output-sensitive directions. Curvature along the output-sensitive directions motivates a candidate monitoring score and a selective regularizer, with an explicit distinction between first-order sensitivity and dynamical instability. The contribution is a technically delimited route from observable preservation to reliability assessment for quantum information-processing components. The numerical study tests only observable preservation and closure size; it does not validate the proposed reliability controller, establish attack detection, or demonstrate improved generalization. No numerical improvement is asserted for the curvature-based mechanism.

Keywords: quantum neural circuits; observable-dependent renormalization; exact model reduction; Krylov closure; quantum Fisher information; Berry curvature; information-system reliability; cybersecurity.

INTRODUCTION

Problem statement. Quantum and hybrid learning systems implement input-dependent computations through parameterized circuits and selected measurements [1], [2]. For an information system, the quantities that must remain correct are usually a finite collection of outputs: class probabilities, expectation values used in a decision rule, or diagnostic measurements. A full description of every density-matrix entry can be unnecessarily expensive for this task. Conversely, a reduction that preserves an average training loss but changes a monitored output is unsuitable as an exact surrogate. The engineering problem is therefore to reduce computational cost while specifying precisely which outputs retain their original dynamics.

This problem is relevant to cybersecurity through the integrity and dependability of quantum information-processing components. Calibration drift, uncontrolled dissipation and parameter perturbations can change a component's outputs. A reduced model can support reference predictions and residual monitoring if it preserves



the quantities actually checked by the surrounding system. It cannot detect changes excluded from its observation contract, and exact simulation of a nominal model is not a guarantee that the nominal model describes a compromised or miscalibrated device. This distinction sets the scope of the reliability claims in this paper.

Analysis of recent research and publications. Near-term quantum algorithms face restrictions on noise, circuit depth and classical optimization [2], [3]. Grigoletto, Tao, Ticozzi and Viola established exact reduction of continuous-time open quantum dynamics using observable Krylov spaces and their extension to operator algebras [4]. Their framework distinguishes a reduced linear model for classical simulation from a physical reduced quantum model. The LA-ODR framework was formulated by Ohnishi, Lemesko and Srichan [5]. Its subsequent application to learnable geometry by Srichan and Ohnishi connects observable-dependent symmetries to learnable quantum geometry and proposes curvature-sensitive control on a reduced space [6]. We use that work for the mathematical definitions needed by the present circuit-reduction problem, rather than pursuing its broader programme.

Trainability and reliability address different questions. Barren plateaus concern vanishing optimization gradients [7], and dynamical Lie algebras provide tools for analysing their occurrence [8]. Quantum Fisher information describes local statistical distinguishability [9], [10] and supports quantum natural-gradient optimization [11]. These connections motivate using geometry after reduction, but neither Lie-algebra dimension nor a curvature value alone establishes prediction accuracy, adversarial robustness or a generalization bound. The unresolved task addressed here is to connect exact output preservation to an operational geometric monitor without silently importing stronger guarantees.

Purpose and contribution. The purpose is to formulate and check a workflow from declared observables to exact reduced dynamics, then specify a proposed reliability mechanism on the resulting model. The contribution is the explicit separation of the algebraic certificate, the numerical observable-preservation check, the conditions for geometric sufficiency, and the unvalidated control proposal. A prospective security application is a nominal reference for output-integrity monitoring of a quantum classifier in an anomaly-detection system, against which measured output probabilities could be compared; attack detection is not evaluated here. The main numerical example remains the structured XX chain and the contrasting generic family; the geometric discussion is restricted to monitoring and selective control of output-sensitive directions.

THEORETICAL BACKGROUND

Quantum circuits with an observation contract

Let the Hilbert space have dimension $N = 2^n$. A parameterized unitary layer has the form $U_\ell(\theta) = \exp(-i \sum_k \theta_k H_k)$, where the generators are Hermitian. On a Markovian continuous-time segment, the density matrix satisfies the Gorini–Kossakowski–Sudarshan–Lindblad equation [12], [13]. Rates are included in the jump operators D_a to keep them distinct from the symmetric logarithmic derivatives used below.

$$\frac{d\rho}{dt} = \mathcal{L}_\theta(t)(\rho) = -i[H_\theta(t), \rho] + \sum_a (D_a \rho D_a^\dagger - \frac{1}{2} \{D_a^\dagger D_a, \rho\}). \quad (1)$$

A fixed time-independent generator defines a semigroup. A driven generator defines a propagator, and a circuit containing arbitrary fixed channels is treated by composition of channels. A general channel is not assumed to be embeddable in a time-independent Lindblad semigroup. The continuous-time numerical example below uses a fixed generator; the closure criterion extends to a family of layers by requiring invariance under every relevant Heisenberg generator or dual channel.

$$\mathcal{O} = \{O_1, \dots, O_m\}, \quad y_i(t; \theta, x) = \text{Tr}[O_i \rho(t; \theta, x)], \quad \rho(0; \theta, x) = \rho_0(x). \quad (2)$$

The target observables are fixed Hermitian operators unless a parameter-dependent family is explicitly declared. Together with the admissible initial states and allowed generators, they form the observation contract. Exactness means equality of every declared output for all allowed initial states and all times in the specified model. Preservation of this contract does not imply reconstruction of the full state. Adding observables can enlarge the required model and can remove a previously available reduction.

Krylov closure and an exact linear realization

For a fixed generator, let $\mathcal{L}^\dagger$ be its adjoint under the trace pairing. The observable Krylov space is

$$\mathcal{K} = \text{span}_{\mathbb{C}} \{(\mathcal{L}^\dagger)^k O_i : i = 1, \dots, m; k \geq 0\}, \quad \mathcal{L}^\dagger \mathcal{K} \subseteq \mathcal{K}. \quad (3)$$

Finite dimensionality bounds its dimension by N^2 . Choose a Hermitian basis $B_1, \dots, B_r$ of any invariant space containing $\mathcal{K}$, with real coefficients A_{ab} and C_{ia} defined by

$$\mathcal{L}^\dagger B_a = \sum_b A_{ab} B_b, \quad O_i = \sum_a C_{ia} B_a, \quad z_a(0) = \text{Tr}[B_a \rho_0]. \quad (4)$$

$$\frac{dz}{dt} = Az, \quad y = Cz. \quad (5)$$

The following proposition states the standard invariant-subspace realization underlying observable



reduction [4]; it is included with its proof to specify the guarantee used here.

Proposition 1. Equations (4)–(5) reproduce the outputs (2) exactly for every initial density matrix. Indeed, differentiating $z_a(t) = \text{Tr}[B_a \rho(t)]$ gives $dz_a/dt = \text{Tr}[(\mathcal{L}^\dagger B_a) \rho(t)] = \sum_\beta A_{a\beta} z_\beta(t)$. Both models therefore solve the same initial-value problem, and the output identity follows. The result is an exact finite-dimensional linear realization; it does not require a positive density matrix interpretation for z .

For piecewise driven or trained circuits, a reusable basis must be invariant under all generators that can occur. Closing only at one parameter value certifies that value. If the basis changes during training, its parameter dependence must be differentiated or a common invariant basis must be constructed. This prevents a pointwise equality from being mistaken for a uniform gradient-preservation theorem.

Pauli-coordinate closure is a convenient, potentially larger realization. Starting with the strings appearing in the target observables, apply the full dual generator to each retained string and add every string with a nonzero coefficient. Repeat until no new string appears. The resulting coordinate span is invariant, but its cardinality need not equal $\dim \mathcal{K}$: a minimal Krylov basis can contain linear combinations whose coefficients are linked by conservation laws. The numerical tables therefore report coordinate-closure sizes rather than minimal dimensions.

Lie symmetries and physical algebraic reduction

For coherent generators, a consistent convention is $\mathfrak{g} = \text{Lie} \mathbb{R}\{iH, k\} \subseteq \mathfrak{u}(N)$. Hermitian and anti-Hermitian parts of other finite-dimensional operators may be included when a larger candidate symmetry algebra is needed. Non-Hermitian jump operators themselves are not elements of $\mathfrak{u}(N)$, and the differential operator $i\partial_t$ is not a finite $N \times N$ matrix. The operational symmetry test is commutation with the entire observable closure:

$$\mathfrak{g}_{\text{ODS}} = \{X \in \mathfrak{g} : [X, B] = 0 \text{ for every } B \in \mathcal{K}\},$$

$$\mathcal{A}_{\text{ODS}} = \{e^X : X \in \mathfrak{g}_{\text{ODS}}\}'. \quad (6)$$

The prime denotes the commutant. Every observable in $\mathcal{K}$ belongs to $\mathcal{A}_{\text{ODS}}$, but a restricted search algebra $\mathfrak{g}$ can miss symmetries and yield an unnecessarily large commutant. Moreover, containment of $\mathcal{K}$ is not itself a proof that this particular algebra is invariant under $\mathcal{L}^\dagger$. A quantum realization requires the operator-algebra construction and its hypotheses, not merely a count of retained Pauli strings [4], [6].

Proposition 2 gives a self-contained sufficient construction within the algebraic reduction approach [4], rather than a claim of a new general reduction theorem. Suppose a unital finite-dimensional $*$ -algebra $\mathcal{A}$ contains the targets and is invariant under $\mathcal{L}^\dagger$. In a suitable unitary basis it has the block form

$$\mathcal{A} \cong \bigoplus_j [M(d_j) \otimes I_{m_j}], \quad \mathcal{H} \cong \bigoplus_j (\mathbb{C}^{d_j} \otimes \mathbb{C}^{m_j}). \quad (7)$$

Write P_j for the block projections and use the ordinary matrix trace on the reduced direct sum. Define compression R by removing inter-block coherences and tracing out each multiplicity factor. Define injection J by restoring the maximally mixed state on that factor:

$$R(\rho) = \bigoplus_j \text{Tr}_{m_j}(P_j \rho P_j), \quad J(\sigma) = \bigoplus_j [\sigma_j \otimes I_{m_j}/m_j]. \quad (8)$$

These maps are completely positive and trace preserving, and RJ is the identity on reduced states. Every $B \in \mathcal{A}$ has the form $\bigoplus_j (B_j \otimes I_{m_j})$, so $\text{Tr}[BJ R(\rho)] = \text{Tr}[B\rho]$. Invariance of $\mathcal{A}$ under $\mathcal{L}^\dagger$ implies that equality of all $\mathcal{A}$ -expectations persists under the full evolution. Consequently $R\mathcal{L} = R\mathcal{L}R$ and the reduced generator $\tilde{\mathcal{L}} = R\mathcal{L}J$ yields $R \exp(t\mathcal{L}) = \exp(t\tilde{\mathcal{L}})R$. Its semigroup is CPTP because $\exp(t\tilde{\mathcal{L}}) = R \exp(t\mathcal{L})J$. This proves exact physical reduction under the stated sufficient hypothesis.

The more general algebraic machinery in [4] addresses the passage from observable spaces to quantum models. Here the explicit construction avoids attributing complete positivity to the linear Pauli-coordinate simulation. Neither Proposition 1 nor Proposition 2 gives a universal polynomial-size model. For an unstructured problem, the observable space or its algebraic extension can fill the original operator space. Depth independence is likewise conditional on a common finite closure and does not imply that closure construction is inexpensive.

RESEARCH METHODOLOGY

Model families and coordinate closure

The numerical study checks the exact linear part of the workflow. Both families use open boundary conditions and the target $O = Z_1$. X_j , Y_j and Z_j denote Pauli operators on site j . In dimensionless units, the structured family S is

$$H_S = \sum_{j=1}^{n-1} (X_j X_{j+1} + Y_j Y_{j+1}) + \sum_{j=1}^n h_j Z_j, \quad D_j = \sqrt{0.3} Z_j. \quad (9)$$

For a fully specified verification, we use $h_j = 0.5 + j/(n+1)$. The contrasting family G is

$$H_G = \sum_{j=1}^{n-1} J_j Z_j Z_{j+1} + \sum_{j=1}^n (a_j X_j + b_j Z_j). \quad (10)$$

Here $J_j = 1 + (j-1)/(2n)$, $a_j = 0.7 + j/(3n)$, $b_j = 0.5 + j/(n+1)$, with dephasing jumps $\sqrt{0.3} Z_j$ and amplitude-damping jumps $\sqrt{0.2} |0\rangle\langle 1|_j$. Fixing these values avoids seed-dependent ambiguity in the support experiment. The amplitude-damping convention is stated explicitly because its adjoint maps Z to a linear combination of Z and the



identity.

Pauli multiplication is implemented through its exact phase table. The support rule is the nonzero structure of each complete dual-generator image; it is not a floating-point Gram–Schmidt rank estimate. All newly generated strings are included, with no truncation of their physical weight. The structured family is examined for $n = 2, \dots, 8$ and the generic family for $n = 2, \dots, 6$. The latter comparison tests this coordinate-closure algorithm, not the impossibility of every alternative reduction.

Independent full and reduced propagation

For $n = 4$ and $n = 5$, construct A from (4) in the retained Pauli basis. Independently evaluate the full density-matrix dynamics; in column-major vectorization, the Hamiltonian part is $-i(I \otimes H - H^T \otimes I)$, and the dephasing term at site j is $0.3(Z_j \otimes Z_j - I)$. The full generator is evaluated directly on density matrices, while the reduced generator acts on Pauli expectations. Both systems are integrated using classical fourth-order Runge–Kutta with step $\Delta t = 0.001$ and sampled at 101 equally spaced times $t_k = 0.05k$, $k = 0, \dots, 100$. A second run at $\Delta t = 0.0005$ checks time-discretization sensitivity.

For each n , a complex matrix B has independent standard-normal real and imaginary entries generated by NumPy default_rng with seed $20261012 + n$. The initial state is $\rho_0 = BB^\dagger / \text{Tr}(BB^\dagger)$. Its Pauli expectations initialize the reduced system, so both propagations start from exactly the same specified state. Double precision is used. The two validation quantities are

$$\begin{aligned} \varepsilon_{\text{out}} &= \max_k |\text{Tr}[Z_1 \rho_{\text{full}}(t_k)] - z_1(t_k)|, \\ \varepsilon_{\text{inv}} &= \max_a \|\mathcal{L}^\dagger B_a - \sum_b A_{ab} B_b\|_F. \end{aligned} \quad (11)$$

The invariance residual checks the embedding of the coordinate model in the full operator dynamics, whereas the output discrepancy checks propagation. Both are numerical diagnostics of the algebraic identity in Proposition 1. Neither is a measurement of hardware noise, classification accuracy, generalization or cybersecurity performance. The implementation uses Python 3.12.14 and NumPy 2.3.5. For a right-hand side F , each Runge–Kutta step uses $k_1 = F(v)$, $k_2 = F(v + \Delta t k_1/2)$, $k_3 = F(v + \Delta t k_2/2)$, $k_4 = F(v + \Delta t k_3)$, followed by $v \leftarrow v + \Delta t(k_1 + 2k_2 + 2k_3 + k_4)/6$. The full and reduced updates use their respective right-hand sides.

Geometry and the scope of preservation

Let $\sigma(q)$ be a smooth family of reduced density matrices on a fixed-rank patch, with coordinates q . Write S_μ for its symmetric logarithmic derivatives, reserving D_a for dissipative jumps. The quantum Fisher matrix F_Q and Bures metric g_B are

$$\partial_\mu \sigma = \frac{1}{2}(S_\mu \sigma + \sigma S_\mu), \quad F_{Q,\mu\nu} = \frac{1}{2} \text{Tr}[\sigma \{S_\mu, S_\nu\}], \quad g_B = F_Q/4. \quad (12)$$

A nondegenerate pullback is a Riemannian metric; a singular pullback requires removal of redundant directions or restriction to a regular identifiable patch [9]–[11]. Preserving a finite list of measurement means does not generally preserve F_Q . The S_μ depend on the state and its derivatives, and the matrix requires their symmetrized products. Merely adding S_μ at one parameter value to the target list is insufficient for a statement over a neighbourhood.

A sufficient geometric condition is a parameter-independent compression with a parameter-independent CPTP recovery that reconstructs the entire state family in the neighbourhood under study. Contractivity in both directions then preserves the Bures distance and its pullback metric. Alternatively, one can explicitly preserve and reconstruct the required state derivatives and SLD products over the declared family. In the absence of such conditions, (12) defines the geometry of the reduced model only; equality with the full-model geometry is not claimed. Computing curvature also requires sufficiently smooth metric derivatives and constant-rank local coordinates.

For a normalized pure-state family, fix the Berry connection convention $A_\mu = i\langle \psi | \partial_\mu \psi \rangle$. The corresponding curvature is

$$F_{\mu\nu} = \partial_\mu A_\nu - \partial_\nu A_\mu = -2\text{Im}\langle \partial_\mu \psi | \partial_\nu \psi \rangle. \quad (13)$$

A small commutator loop generated by Hermitian HA and HB has leading correction $-\alpha\beta[HA, HB]$. This explains why noncommutativity can contribute to geometric phase, but a nonzero operator commutator need not have a nonzero expectation in the particular state. Berry curvature is distinct from the Riemann curvature of g_B . Mixed-state phase constructions require additional choices and are not inferred from (13) by replacing a vector with a density matrix. The monitoring proposal below uses Riemannian curvature.

RESEARCH RESULTS

Exactness and dimension of the observable realization

Table 1 reports the Pauli-coordinate closures obtained by the procedure in Section 3. The structured count has a direct algebraic explanation. Jordan–Wigner Majorana operators $c_1, \dots, c_{2n}$ satisfy canonical anticommutation relations. Their bilinears $ic_a c_b$, $a < b$, span a space of dimension $\binom{2n}{2} = n(2n-1)$. Commutation with a quadratic Hamiltonian preserves this space, and conjugation by a local Z_j changes only the sign of a Pauli string.



Thus local dephasing preserves the bilinear span as well. Since Z_1 is a bilinear, the Heisenberg trajectory remains in this invariant space.

Table 1

Pauli-coordinate closure sizes for the specified models

n	Full space 4^n	Model S	$n(2n-1)$	Model G
2	16	6	6	16
3	64	15	15	64
4	256	28	28	256
5	1024	45	45	1024
6	4096	66	66	4096
7	16384	91	91	Not evaluated
8	65536	120	120	Not evaluated

For the specified nonzero fields, Pauli-coordinate closure reaches this bilinear support. The count is a quadratic upper bound for the minimal linear realization, not a proof of minimality. In particular, the XX Hamiltonian conserves excitation number and admits additional relations between Majorana coefficients. The span of centred number-conserving fermion quadratics has dimension n^2 and is also invariant for this model. A minimal single-observable Krylov space can therefore be smaller than the Pauli support reported in Table 1. This does not impair exactness of the larger coordinate realization.

For five qubits, the coordinate count is reduced by a factor of $1024/45 \approx 22.8$ relative to the full operator space. At eight qubits it is $65536/120 \approx 546.1$. These are representation-size ratios; they are not measured runtime or memory speedups. They also do not describe the dimension of an associative reduced quantum algebra. For model G, the coordinate closure fills the full Pauli basis at every tested size. This supplies a concrete failure case for assuming that a small target set alone guarantees compression.

Numerical check of observable preservation

Table 2 compares the two independently constructed propagations for model S. The invariance residual and the trajectory discrepancies are consistent with double-precision evaluation of an exactly invariant realization. The check concerns a fixed initial state and a finite time grid for each size; exactness for all states and times follows from Proposition 1 and the invariant-space construction, rather than from extrapolating these samples.

Table 2

Numerical verification of model S on $t \in [0,5]$

n	Coordinates	ε_{out}	ε_{inv}	Step-halving change
4	28 / 256	4.79e-16	1.79e-15	6.15e-13
5	45 / 1024	4.40e-16	1.99e-15	3.34e-13

Halving the integration step changes the full-model output by 6.15×10^{-13} at $n = 4$ and 3.34×10^{-13} at $n = 5$. This is a convergence diagnostic, not a rigorous bound on time-discretization error. The much smaller full-versus-reduced discrepancy measures consistency between the two realizations under the same integrator; it must not be interpreted as the absolute error against the exact continuous-time solution. The experiment does not test the CPTP maps or the proposed geometric controller.

A practical failure mode is unstable numerical rank estimation. Repeatedly applying a generator can amplify tiny round-off components outside an intended invariant subspace. A stopping rule based solely on a tolerance-dependent Gram–Schmidt rank can then either overestimate the space or discard legitimate directions. Structural Pauli closure and direct invariance checks avoid that particular ambiguity for the present family. They do not eliminate coefficient round-off during propagation, and they do not make a large generic closure computationally cheap.

Operational protected and dangerous directions

The proposed reliability extension follows the orbit/protected/dangerous distinction in [6], but makes its local meaning explicit. On a regular reduced model manifold X , let $f(p)$ be the vector of declared outputs and choose a positive metric g . Suppose an effective symmetry action leaves f invariant. Its orbit tangent space Torb



lies in $\ker df_p$. Define the protected normal space and its orthogonal complement by

$$N_{\text{prot}} = \ker df_p \cap T_{\text{orb}}^\perp, \quad N_{\text{dang}} = (\ker df_p)^\perp. \quad (14)$$

$$T_p X = T_{\text{orb}} \oplus N_{\text{prot}} \oplus N_{\text{dang}}. \quad (15)$$

This is a pointwise decomposition; smooth projectors require constant rank and a fixed orbit type locally. “Dangerous” means output-sensitive to first order, not necessarily unstable or malicious. For a small tangent displacement v , $df_p(v) = df_p(\Pi_{\text{dang}} v)$. Consequently the declared output variation obeys a local first-order estimate

$$\|f(\exp_p v) - f(p)\| \leq \|df_p\| \|\Pi_{\text{dang}} v\|_g + O(\|v\|_g^2). \quad (16)$$

Equation (16) follows by a Taylor expansion in normal coordinates. It is a local output-sensitivity statement, not a statistical generalization bound. A vector in $\ker df_p$ can affect an output at second order or become visible after evolution. Exact dynamical protection requires preservation of the relevant output trajectory, equivalently the full observable closure in the linear setting. A monitor should therefore use trajectory-sensitive outputs or separately test invariance of a proposed protected distribution.

The symmetry group defining a commutant needs special care. Conjugation by GODS acts trivially on its own commutant $\mathcal{A}ODS$. Hence a nontrivial residual orbit on a reduced parameter or state manifold cannot be inferred simply by choosing $K \subset GODS$. Any effective action used in (14) must be specified independently, for example as a parameter redundancy on a larger representation. When no such action is established, set $T_{\text{orb}} = \{0\}$; the output-Jacobian decomposition remains meaningful without a Hamiltonian normal-form claim.

Centre-manifold and normal-form methods may be useful for an additional nonlinear learning dynamics, but their spectral and regularity hypotheses must be checked [14]. A general dissipative state space does not automatically carry a Hamiltonian compact-group action. Nor does a spectral gap alone make every proposed centre manifold attracting when unstable modes are present. We therefore do not use a normal-form theorem as a universal certificate for (15), and we do not identify the word “protected” with nonlinear stability.

Proposed curvature-based monitoring and control

Within a regular geodesically convex neighbourhood, a reference trajectory $\bar{\gamma}(t)$ and a monitored trajectory $\gamma(t)$ can be compared by the tangent displacement $v(t) = \text{Log} \bar{\gamma}(t) \gamma(t)$. Subtracting manifold points as if they were vectors is avoided. For a unit reference tangent u , define a local adverse-curvature score using the sign convention in which the Jacobi equation is $\nabla u^2 J + R(J, u)u = 0$:

$$\kappa_-(p, u) = \max(0, -\min_w g(R(w, u)u, w)), \quad w \in N_{\text{dang}}, \quad \|w\|_g = 1, \quad w \perp u. \quad (17)$$

If the admissible set of w is empty, set $\kappa_- = 0$. Negative sectional curvature permits separation of nearby geodesics under this convention, which motivates the score; the actual learning or device trajectory need not be geodesic. Positive curvature, parameter forcing and non-normal dynamics can also affect reliability. Equation (17) is therefore a diagnostic feature rather than a necessary or sufficient instability test.

$$S(t) = \|\Pi_{\text{dang}} v(t)\|_g^2 + \beta \kappa_-(\bar{\gamma}(t), u(t)), \quad \beta \geq 0. \quad (18)$$

Both terms in (18) are nonnegative. β sets their relative scale and units, and must be calibrated along with any alarm threshold. This operational form adapts the geometric penalty proposed in [6]; it does not interpret signed Ricci curvature as a calibrated probability of failure. At a stationary reference point, the curvature term can be omitted or evaluated over specified candidate directions. When the reference leaves a regular convex neighbourhood, the local construction must be rebuilt rather than continued through a singular chart.

A complementary candidate regularizer can be defined if reduced Hermitian layer generators $\tilde{H}\ell$ exist. Set $C\ell = i[\tilde{H}\ell, \tilde{H}\ell_{+1}]$, which is Hermitian, and map it to the state tangent $V\ell = -i[C\ell, \sigma]$. Then

$$\mathcal{L}_{\text{train}} = \mathcal{L}_{\text{task}} + \lambda \sum_\ell \|\Pi_{\text{dang}} V\ell\|_g^2, \quad \lambda \geq 0. \quad (19)$$

This specifies what is projected: a tangent vector, not an operator implicitly identified with a tangent vector. It avoids penalizing protected commutator-induced motion at first order. It is not asserted that (19) minimizes (17), improves prediction accuracy or preserves every nonlinear representation property. If only a linear coordinate realization is available, there need not be physical reduced Hamiltonians $\tilde{H}\ell$, and this particular regularizer is not automatically implementable.

For an information-processing component, a possible control loop evaluates the declared outputs and local score, logs deviations from a reference envelope, and tests a bounded adjustment of circuit parameters or invokes a full-model check. Every candidate adjustment must remain within the certified generator family; otherwise closure must be recomputed. Thresholds, actuation limits and the reference envelope require a validation set representing the intended device and perturbations. The numerical study in this paper does not tune them.

Evidence boundaries and cybersecurity interpretation

Table 3 summarizes the evidence attached to each step. The algebraic propositions establish output preservation within a declared mathematical model. The numerical experiment tests an implementation of that result. The geometric stage introduces additional state-family, smoothness and observability conditions. The



control stage remains a proposal. This hierarchy prevents numerical precision in a simulator from being presented as evidence of operational robustness.

Table 3

Scope and status of the claims

Construction	Evidence	Scope
Invariant linear reduction	Proposition 1 and Table 2	Exact declared outputs within the specified dynamics
CPTP quantum realization	Proposition 2; reference [4]	Requires the stated algebraic hypothesis; not tested in Table 2
Geometric preservation	Sufficient conditions in Section 3.3	Not implied by preservation of a few measurement means
Protected/dangerous split	Local output-Jacobian construction	First-order output sensitivity; no automatic stability guarantee
Curvature monitor and regularizer	Proposed equations (17)–(19)	No numerical reliability improvement or attack-detection result

For cybersecurity, the useful output of exact reduction is a traceable nominal behaviour model. An integrity check can compare measured output values with that model, subject to uncertainty in state preparation, noise calibration and measurement statistics. The proposed directional score could help identify which local parameter changes affect declared outputs. However, neither a small residual nor a small curvature score excludes an attack on an unobserved quantity. No attack dataset, adversarial optimization experiment, hardware trial or detection-rate evaluation is reported here.

Generalization also requires separate evidence. Bounds on local curvature and Lipschitz constants do not alone control the capacity of a trained hypothesis class or its sampling error. For this reason, the curvature-dependent generalization expression sketched in [6] is not promoted to a theorem in this article. The validated claim is exact reduction for the specified observables; the proposed claim is that directional geometric information may support monitoring and selective intervention. No numerical improvement is asserted for that proposal.

The principal limitations are finite-dimensional Markovian modelling, possible exponential closure growth, the distinction between coordinate and quantum-algebra dimensions, and the cost of preserving geometry rather than a few readouts. The examples concern one target observable and two circuit-generator families. A practical implementation also has to account for changes in the observable set, coordinate conditioning, derivative estimation and noise-model mismatch. These restrictions define where the reduction certificate applies and where an empirical reliability evaluation is still required.

CONCLUSIONS AND PROSPECTS FOR FURTHER RESEARCH

An observable contract provides a precise starting point for reducing quantum neural and information-processing circuits. Closure under the dual dynamics yields an exact linear realization of the declared outputs, while a physically admissible quantum realization requires an additional operator-algebra construction. The explicit sufficient CPTP construction and the coordinate-space proof state these two claims separately. A reduced dimension is determined jointly by the dynamics and the targets, and can equal the unreduced dimension.

For the structured XX chain with local dephasing, the Pauli-coordinate realization has $n(2n-1)$ coordinates and reproduces the selected magnetization output within numerical precision in the specified four- and five-qubit checks. The generic comparison fills the Pauli basis at the tested sizes. The quadratic count is an invariant representation size, not a minimal Krylov dimension, a universal compression guarantee, or a demonstrated hardware speedup.

Reduced quantum geometry can support a local decomposition into gauge, output-insensitive and output-sensitive directions when its sufficiency and regularity conditions are satisfied. A nonnegative directional monitoring score and a tangent-space commutator regularizer are proposed for the output-sensitive block. Their reliability benefit remains unvalidated. Future experiments should compare unregularized, full-commutator and directional-regularizer baselines under identical resources; report output error, false-alarm rate, detection delay and intervention cost; and test whether signals remain informative under model mismatch. Such experiments are needed before claiming cybersecurity detection or improved deployed reliability.



CODE AND DATA AVAILABILITY

Appendix A provides the complete Python script used to generate Tables 1 and 2, with execution instructions and reference outputs. The same script is supplied as a machine-readable file in the accompanying supplementary archive, which contains `reproduce_results.py`, the recorded `verification_results.json`, `requirements.txt` and a `README`. The script generates the closure dimensions in Table 1 and the trajectory, invariance and step-halving diagnostics in Table 2 without external datasets. All model parameters and random seeds are specified in Section 3 and in the script. Results were generated with Python 3.12.14 and NumPy 2.3.5; floating-point residuals near machine precision may vary across numerical environments. The supplement does not implement or validate the proposed geometric controller.

USE OF ARTIFICIAL INTELLIGENCE TOOLS

The earlier journal draft records the use of Claude (Anthropic) for drafting and its preliminary numerical illustration. Codex (OpenAI) was used in preparing the present revision for language editing, mathematical consistency checks, numerical verification code and document formatting. AI tools are not authors. Responsibility for the scientific claims, author information and submission rests with the human authors.

REFERENCES

1. Biamonte, J., Wittek, P., Pancotti, N., Rebentrost, P., Wiebe, N., & Lloyd, S. (2017). Quantum machine learning. *Nature*, 549, 195–202. <https://doi.org/10.1038/nature23474>
2. Cerezo, M., Arrasmith, A., Babbush, R., Benjamin, S. C., Endo, S., Fujii, K., McClean, J. R., Mitarai, K., Yuan, X., Cincio, L., & Coles, P. J. (2021). Variational quantum algorithms. *Nature Reviews Physics*, 3, 625–644. <https://doi.org/10.1038/s42254-021-00348-9>
3. Preskill, J. (2018). Quantum computing in the NISQ era and beyond. *Quantum*, 2, 79. <https://doi.org/10.22331/q-2018-08-06-79>
4. Grigoletto, T., Tao, Y., Ticozzi, F., & Viola, L. (2025). Exact model reduction for continuous-time open quantum dynamics. *Quantum*, 9, 1814. <https://doi.org/10.22331/q-2025-07-29-1814>
5. Ohnishi, I., Lemesko, A., & Srichan, C. (2026). Lie-algebraic observable-dependent renormalization: A unified framework for driven open quantum systems and stochastic Navier–Stokes equations [Preprint]. SSRN. <https://doi.org/10.2139/ssrn.6681038>
6. Srichan, C., & Ohnishi, I. (2026). Quantum learnable geometry and Lie-algebraic observable-dependent renormalization (Rev. 01) [Preprint]. SSRN. <https://doi.org/10.2139/ssrn.7444558>
7. McClean, J. R., Boixo, S., Smelyanskiy, V. N., Babbush, R., & Neven, H. (2018). Barren plateaus in quantum neural network training landscapes. *Nature Communications*, 9, 4812. <https://doi.org/10.1038/s41467-018-07090-4>
8. Larocca, M., Czarnek, P., Sharma, K., Muraleedharan, G., Coles, P. J., & Cerezo, M. (2022). Diagnosing barren plateaus with tools from quantum optimal control. *Quantum*, 6, 824. <https://doi.org/10.22331/q-2022-09-29-824>
9. Braunstein, S. L., & Caves, C. M. (1994). Statistical distance and the geometry of quantum states. *Physical Review Letters*, 72, 3439–3443. <https://doi.org/10.1103/PhysRevLett.72.3439>
10. Meyer, J. J. (2021). Fisher information in noisy intermediate-scale quantum applications. *Quantum*, 5, 539. <https://doi.org/10.22331/q-2021-09-09-539>
11. Stokes, J., Isaac, J., Killoran, N., & Carleo, G. (2020). Quantum natural gradient. *Quantum*, 4, 269. <https://doi.org/10.22331/q-2020-05-25-269>
12. Gorini, V., Kossakowski, A., & Sudarshan, E. C. G. (1976). Completely positive dynamical semigroups of N-level systems. *Journal of Mathematical Physics*, 17, 821–825. <https://doi.org/10.1063/1.522979>
13. Lindblad, G. (1976). On the generators of quantum dynamical semigroups. *Communications in Mathematical Physics*, 48, 119–130. <https://doi.org/10.1007/BF01608499>
14. Hochs, P., & Roberts, A. J. (2019). Normal forms and invariant manifolds for nonlinear, non-autonomous PDEs, viewed as ODEs in infinite dimensions. *Journal of Differential Equations*, 267, 7263–7312. <https://doi.org/10.1016/j.jde.2019.07.021>

**Чавіс Срічан**

кафедра комп'ютерної інженерії,
Університет Кхонкена, Кхонкен, Таїланд
ORCID: 0000-0002-4772-3543
chavis@kku.ac.th

Ісаму Онісі

Вища школа інтегрованих наук про життя,
Університет Хіросіми, Хігасі-Хіросіма, Японія
ORCID: 0009-0009-5716-6216
isamu_o@toki.waseda.jp

ТОЧНА АЛГЕБРАІЧНА РЕДУКЦІЯ КВАНТОВИХ НЕЙРОННИХ СХЕМ: ЗАЛЕЖНА ВІД СПОСТЕРЕЖУВАНИХ РЕНОРМАЛІЗАЦІЯ ТА КЕРУВАННЯ НАДІЙНІСТЮ НА ОСНОВІ КРИВИНИ

Анотація. Обчислювальна складність і надійність параметризованих квантових схем залежать від того, які вихідні величини має зберігати інформаційна система. У статті розроблено послідовність побудов від квантової нейронної схеми та заданих спостережуваних через замикання Лі й Крилова до точної алгебраїчної редукції та запропонованого геометричного моніторингу надійності. Динаміку описано скінченновимірною марковською моделлю з явно визначеним набором цільових спостережуваних. Розмежовано мінімальний лінійний простір спостережуваних, замкнений координатний простір операторів Паулі для класичного моделювання та асоціативну операторну алгебру, на якій може бути реалізована цілком додатна редуктована квантова динаміка зі збереженням сліду. Інваріантність простору дає безпосередній доказ точного відтворення заданих виходів; окремо наведено достатню конструкцію фізично допустимого квантового скорочення. Структурований ХХ-ланцюжок із локальним дефазуванням ілюструє перехід від повного операторного простору розмірності 4^n до простору з $n(2n-1)$ координатами Паулі, зокрема від 1024 до 45 координат для п'яти кубітів. Незалежна чисельна перевірка з фіксованими параметрами порівнює повну та редуктовану еволюцію, а модель із загальнішими взаємодіями демонструє відсутність скорочення тим самим методом координатного замикання. Квантову інформацію Фішера та кривину Беррі визначено з додатковими умовами достатності й гладкості; їх збереження не виводиться лише зі збереження середніх значень кількох вимірювань. Локальний розклад розділяє калібрувальні, нечутливі до виходів і чутливі до виходів напрямки. Для останніх запропоновано показник моніторингу на основі кривини та вибіркового регуляризатор. Чутливість першого порядку відокремлено від динамічної нестійкості. Практична спрямованість полягає у формуванні перевірюваної номінальної моделі квантового компонента інформаційної системи. Чисельні результати стосуються лише точності відтворення спостережуваних і розміру замикання. Поліпшення надійності, узагальнювальної здатності або виявлення кібератак за допомогою запропонованого механізму чисельно не підтверджено; чисельне поліпшення для механізму керування на основі кривини не заявляється.

Ключові слова: квантові нейронні схеми; залежна від спостережуваних ренормалізація; точна редукція моделей; замикання Крилова; квантова інформація Фішера; кривина Беррі; надійність інформаційних систем; кібербезпека.

СПИСОК ВИКОРИСТАНИХ ДЖЕРЕЛ

1. Biamonte, J., Wittek, P., Pancotti, N., Rebentrost, P., Wiebe, N., & Lloyd, S. (2017). Quantum machine learning. *Nature*, 549, 195–202. <https://doi.org/10.1038/nature23474>
2. Cerezo, M., Arrasmith, A., Babbush, R., Benjamin, S. C., Endo, S., Fujii, K., McClean, J. R., Mitarai, K., Yuan, X., Cincio, L., & Coles, P. J. (2021). Variational quantum algorithms. *Nature Reviews Physics*, 3, 625–644. <https://doi.org/10.1038/s42254-021-00348-9>
3. Preskill, J. (2018). Quantum computing in the NISQ era and beyond. *Quantum*, 2, 79. <https://doi.org/10.22331/q-2018-08-06-79>
4. Grigoletto, T., Tao, Y., Ticozzi, F., & Viola, L. (2025). Exact model reduction for continuous-time open



- quantum dynamics. *Quantum*, 9, 1814. <https://doi.org/10.22331/q-2025-07-29-1814>
5. Ohnishi, I., Lemesko, A., & Srichan, C. (2026). Lie-algebraic observable-dependent renormalization: A unified framework for driven open quantum systems and stochastic Navier–Stokes equations [Preprint]. SSRN. <https://doi.org/10.2139/ssrn.6681038>
 6. Srichan, C., & Ohnishi, I. (2026). Quantum learnable geometry and Lie-algebraic observable-dependent renormalization (Rev. 01) [Preprint]. SSRN. <https://doi.org/10.2139/ssrn.7444558>
 7. McClean, J. R., Boixo, S., Smelyanskiy, V. N., Babbush, R., & Neven, H. (2018). Barren plateaus in quantum neural network training landscapes. *Nature Communications*, 9, 4812. <https://doi.org/10.1038/s41467-018-07090-4>
 8. Larocca, M., Czarnik, P., Sharma, K., Muraleedharan, G., Coles, P. J., & Cerezo, M. (2022). Diagnosing barren plateaus with tools from quantum optimal control. *Quantum*, 6, 824. <https://doi.org/10.22331/q-2022-09-29-824>
 9. Braunstein, S. L., & Caves, C. M. (1994). Statistical distance and the geometry of quantum states. *Physical Review Letters*, 72, 3439–3443. <https://doi.org/10.1103/PhysRevLett.72.3439>
 10. Meyer, J. J. (2021). Fisher information in noisy intermediate-scale quantum applications. *Quantum*, 5, 539. <https://doi.org/10.22331/q-2021-09-09-539>
 11. Stokes, J., Izaac, J., Killoran, N., & Carleo, G. (2020). Quantum natural gradient. *Quantum*, 4, 269. <https://doi.org/10.22331/q-2020-05-25-269>
 12. Gorini, V., Kossakowski, A., & Sudarshan, E. C. G. (1976). Completely positive dynamical semigroups of N-level systems. *Journal of Mathematical Physics*, 17, 821–825. <https://doi.org/10.1063/1.522979>
 13. Lindblad, G. (1976). On the generators of quantum dynamical semigroups. *Communications in Mathematical Physics*, 48, 119–130. <https://doi.org/10.1007/BF01608499>
 14. Hochs, P., & Roberts, A. J. (2019). Normal forms and invariant manifolds for nonlinear, non-autonomous PDEs, viewed as ODEs in infinite dimensions. *Journal of Differential Equations*, 267, 7263–7312. <https://doi.org/10.1016/j.jde.2019.07.021>

Отримано редакцією журналу / Received: 22.08.26

Прорецензовано / Revised: 30.08.26

Схвалено до друку / Accepted: 26.10.26



This work is licensed under Creative Commons Attribution-noncommercial-sharealike 4.0 International License.