Coherent Lyapunov Feedback Control for Quantum Machine Learning

A Novel Algorithm for Parameter Optimization in Parameterized Quantum Circuits

Mark Luo

SuperQubit Inc.
Markham, ON, Canada

June 2026

Abstract

We propose Coherent Lyapunov Feedback (CLF), a fundamentally new training paradigm for parameterized quantum circuits (PQCs) that eliminates projective measurement entirely from the parameter optimization loop. Classical quantum machine learning (QML) methods, including the parameter-shift rule and its generalizations, require $\mathcal{O}(M)$ circuit evaluations and $\mathcal{O}(M \times S)$ total measurement shots per gradient step, where M is the number of parameters and S is the number of shots per evaluation. CLF replaces this with a closed-loop quantum feedback architecture: a target state $|\psi_{\text{target}}\rangle$ acts as a reference, a coherent overlap module computes the complex fidelity amplitude between the current output and the target without wavefunction collapse, and a Lyapunov feedback law drives continuous parameter updates with mathematically guaranteed monotonic convergence. We formalize the coherent overlap circuit, derive the Lyapunov stability proof for the PQC setting, establish the fixed-point topology, analyze noise robustness under Lindblad dynamics, and present the complete CLF algorithm with complexity analysis. CLF reduces measurement overhead to $\mathcal{O}(S)$ ancilla measurements independent of the number of parameters M while preserving global convergence guarantees absent in gradient-based methods. The framework opens a new design space at the intersection of quantum control theory and quantum machine learning.

Keywords: quantum machine learning, parameterized quantum circuits, Lyapunov stability, coherent feedback control, measurement-free optimization, quantum fidelity, barren plateaus, variational quantum algorithms

Contents

1 Introduction

Parameterized quantum circuits (PQCs) sit at the heart of near-term quantum machine learning. A PQC encodes a task as the optimization of a loss function

$$L(\boldsymbol{\theta}) = \langle\psi_0|\, U(\boldsymbol{\theta})^\dagger \hat{O}\, U(\boldsymbol{\theta})\, |\psi_0\rangle \,,$$

where $U(\boldsymbol{\theta})$ is a product of parameterized gates and $\hat{O}$ is an observable encoding the objective. The dominant training method, the parameter-shift rule, computes the exact gradient via:

$$\frac{\partial L}{\partial \theta_k} = \frac{L\big(\theta_k + \frac{\pi}{2}\big) - L\big(\theta_k - \frac{\pi}{2}\big)}{2}.$$

This requires $2M$ separate circuit executions per gradient step, each demanding $S \gg 1$ measurement shots to estimate expectation values reliably. For $M = 100$ parameters and $S = 1000$ shots, a single gradient step costs 200,000 circuit evaluations. This measurement bottleneck is a primary obstacle to scaling QML on near-term hardware.

Beyond efficiency, gradient-based methods suffer from the *barren plateau* phenomenon: the gradient magnitude decreases exponentially with system size, making optimization increasingly ineffective as circuits grow. These two problems—measurement cost and barren plateaus—motivate the search for fundamentally different training paradigms.

In this paper, we introduce Coherent Lyapunov Feedback (CLF), a training algorithm that eliminates projective measurement from the optimization loop. The key insight is to reframe QML training as a quantum feedback control problem: rather than computing a loss and descending its gradient, we set a target state $|\psi_{\text{target}}\rangle$ as a reference signal and design a feedback law that continuously drives the circuit output $U(\boldsymbol{\theta})\,|\psi_0\rangle$ toward $|\psi_{\text{target}}\rangle$ using only coherent quantum operations. The stability of convergence is guaranteed by Lyapunov's second method, a classical framework from control theory that we here formalize for the PQC setting.

1.1 Summary of Contributions

- **Formalization of the Coherent Overlap Module (COM)**: a quantum circuit that computes the complex fidelity amplitude $\langle\psi_{\text{target}}|U(\boldsymbol{\theta})|\psi_0\rangle$ as a coherent phase, without collapsing either state.
- **Lyapunov stability proof for PQCs**: rigorous derivation of the CLF parameter update law, proof of monotonic convergence, and characterization of fixed-point topology.
- **Noise robustness analysis**: behavior of CLF under Lindblad master equation dynamics, including conditions for convergence under decoherence.
- **Complete CLF algorithm specification** with circuit construction, complexity analysis, and comparison to parameter-shift and gradient-free baselines.
- **Open problem formulation**: connections to compressed sensing on the loss landscape and Lie-algebraic circuit design for efficient CLF deployment.

2 Background and Problem Formulation

2.1 Parameterized Quantum Circuits

Let $\mathcal{H} = (\mathbb{C}^2)^{\otimes n}$ be an n-qubit Hilbert space. A PQC is a unitary operator:

$$U(\boldsymbol{\theta}) = \prod_{k=1}^{M} \exp(-i\,\theta_k G_k)\,,$$

where $\boldsymbol{\theta} = (\theta_1, \ldots, \theta_M) \in \mathbb{R}^M$ is the parameter vector and $\{G_k\}$ are Hermitian generators (typically tensor products of Pauli matrices). Given an initial state $|\psi_0\rangle$ and a target state $|\psi_{\text{target}}\rangle$, the training objective is:

$$\text{maximize} \quad F(\boldsymbol{\theta}) = |\langle\psi_{\text{target}}|U(\boldsymbol{\theta})|\psi_0\rangle|^2 \in [0, 1],$$

with $F = 1$ achieved if and only if $U(\boldsymbol{\theta})\,|\psi_0\rangle = e^{i\phi}\,|\psi_{\text{target}}\rangle$ for some global phase ϕ. This formulation covers state preparation, quantum compilation, unitary learning, and classification tasks where $|\psi_{\text{target}}\rangle$ encodes the label.

2.2 The Measurement Cost Problem

The parameter-shift rule gives the exact gradient component:

$$\frac{\partial F}{\partial \theta_k} = F\big(\theta_k + \tfrac{\pi}{2}\big) - F\big(\theta_k - \tfrac{\pi}{2}\big)\,.$$

Each evaluation $F(\boldsymbol{\theta})$ requires S measurement shots on the observable $|\psi_{\text{target}}\rangle\,\langle\psi_{\text{target}}|$. The total shot cost per gradient step scales as:

$$\text{Cost}_{\text{PS}} = 2M \cdot S \text{ shots.}$$

For realistic circuits ($M \sim 10^2$ to 10^3, $S \sim 10^3$ to 10^4), this is 10^5 to 10^7 shots per step, with potentially thousands of steps required for convergence. On near-term hardware with shot rates of $\sim\!10^4$ per second, this translates to hours or days of runtime.

2.3 Lyapunov Stability Theory (Classical)

For a dynamical system $\dot{x} = f(x)$, Lyapunov's second method establishes asymptotic stability of an equilibrium x^* by finding a scalar function $V(x)$ satisfying:

$$V(x) > 0 \text{ for all } x \neq x^*, \qquad V(x^*) = 0, \qquad \text{and} \qquad \dot{V} \leq 0 \text{ along trajectories.}$$

If $\dot{V} < 0$ everywhere except at x^*, convergence to x^* is guaranteed regardless of initial conditions. We will use this framework with $V(\boldsymbol{\theta}) = 1 - F(\boldsymbol{\theta})$ as the Lyapunov function, exploiting the fact that $F \in [0, 1]$ ensures $V \geq 0$ with $V = 0$ iff $\boldsymbol{\theta}$ is optimal.

3 The Coherent Overlap Module

The fundamental challenge in measurement-free feedback is: how can the circuit parameters "feel" the discrepancy between the current state and the target without a projective measurement collapsing the quantum state? We answer this by constructing the *Coherent Overlap Module* (COM), a quantum circuit that encodes the complex overlap amplitude:

$$A(\boldsymbol{\theta}) := \langle \psi_{\text{target}} | U(\boldsymbol{\theta}) | \psi_0 \rangle \in \mathbb{C}$$

into the phase of an ancilla qubit through interference, without measuring or disturbing either $|\psi(\boldsymbol{\theta})\rangle = U(\boldsymbol{\theta}) |\psi_0\rangle$ or $|\psi_{\text{target}}\rangle$.

3.1 The Generalised Hadamard Test for Two Distinct States

The standard Hadamard test estimates $\text{Re}[\langle\psi| V |\psi\rangle]$ for a unitary V acting on a *single* state $|\psi\rangle$. To extract the complex amplitude $A(\boldsymbol{\theta}) = \langle B|A\rangle$ between two *distinct* states $|A\rangle = U(\boldsymbol{\theta}) |\psi_0\rangle$ and $|B\rangle = |\psi_{\text{target}}\rangle$, we use a generalised two-state Hadamard test. The circuit uses two registers: one ancilla qubit and one n-qubit work register initialised to $|\psi_0\rangle$. The controlled operations prepare $|A\rangle$ or $|B\rangle$ in the work register conditioned on the ancilla, creating interference that encodes $A(\boldsymbol{\theta})$ as a coherent phase.

Construction 1 (Coherent Overlap Module (COM) — Generalised Hadamard Test)**.** *Registers:* $|\text{anc}\rangle = |0\rangle$ *(1 qubit), work register* $|w\rangle = |\psi_0\rangle$ *(n qubits). No register B is required separately;* $|\psi_{\text{target}}\rangle$ *is encoded in the controlled preparation* V_B *below.*

Define two n-qubit unitaries acting on the work register:

$$V_A := U(\boldsymbol{\theta}), \quad \textit{so that } V_A |\psi_0\rangle = |A\rangle \,,$$
$$V_B := \textit{any unitary such that } V_B |\psi_0\rangle = |B\rangle = |\psi_{target}\rangle \,.$$

Step 1. *Apply H to the ancilla:*

$$|\Psi_1\rangle = \frac{1}{\sqrt{2}}\big(|0\rangle + |1\rangle\big) \otimes |\psi_0\rangle \,.$$

Step 2. *Apply* $V_B^\dagger$ *controlled on ancilla* $= |0\rangle$*, and* V_A *controlled on ancilla* $= |1\rangle$*:*

$$|\Psi_2\rangle = \frac{1}{\sqrt{2}}\Big(|0\rangle \otimes V_B^\dagger |\psi_0\rangle \;+\; |1\rangle \otimes V_A |\psi_0\rangle\Big) = \frac{1}{\sqrt{2}}\Big(|0\rangle |B^*\rangle + |1\rangle |A\rangle\Big),$$

where $|B^*\rangle := V_B^\dagger |\psi_0\rangle$*. Since* $V_B |\psi_0\rangle = |B\rangle$*, we have* $|B^*\rangle = V_B^\dagger |\psi_0\rangle$*, so* $\langle\psi_0|V_B|B^*\rangle = \langle\psi_0|\psi_0\rangle = 1$ *and* $\langle A|B^*\rangle = \langle\psi_0| V_A^\dagger V_B^\dagger |\psi_0\rangle$*.*

Simpler equivalent: *initialise work register to* $|0\rangle^{\otimes n}$*, apply* $\widetilde{V}_B^\dagger$ *(ctrl-0) and* $\widetilde{V}_A$ *(ctrl-1) where* $\widetilde{V}_X |0\rangle^{\otimes n} = |X\rangle$*. Then Step 2 gives exactly:*

$$|\Psi_2\rangle = \frac{1}{\sqrt{2}}\Big(|0\rangle |B\rangle + |1\rangle |A\rangle\Big).$$

Step 3. *Apply H to the ancilla. Using* $H|0\rangle = \frac{1}{\sqrt{2}}(|0\rangle + |1\rangle)$ *and* $H|1\rangle =$

$\frac{1}{\sqrt{2}}(|0\rangle - |1\rangle)$:

$$\begin{aligned}|\Phi\rangle &= (H \otimes I)\, |\Psi_2\rangle \\ &= \frac{1}{2}\Big[(|0\rangle + |1\rangle)\,|B\rangle + (|0\rangle - |1\rangle)\,|A\rangle\Big] \\ &= \frac{1}{2}\Big[|0\rangle\,\big(|B\rangle + |A\rangle\big) + |1\rangle\,\big(|B\rangle - |A\rangle\big)\Big].\end{aligned}$$

Reduced density matrix of the ancilla. *Trace out the work register from $|\Phi\rangle\langle\Phi|$. Writing $|\Phi\rangle = |0\rangle|\phi_+\rangle + |1\rangle|\phi_-\rangle$ with $|\phi_\pm\rangle = \frac{1}{2}(|B\rangle \pm |A\rangle)$, the matrix elements of $\rho_{\text{anc}} = \text{Tr}_w[|\Phi\rangle\langle\Phi|]$ are:*

$$\begin{aligned}\rho_{00} &= \langle\phi_+|\phi_+\rangle = \frac{1}{4}\big(\langle B|B\rangle + \langle B|A\rangle + \langle A|B\rangle + \langle A|A\rangle\big) = \frac{1}{2}\big(1 + \text{Re}\langle A|B\rangle\big), \\ \rho_{11} &= \langle\phi_-|\phi_-\rangle = \frac{1}{4}\big(\langle B|B\rangle - \langle B|A\rangle - \langle A|B\rangle + \langle A|A\rangle\big) = \frac{1}{2}\big(1 - \text{Re}\langle A|B\rangle\big), \\ \rho_{01} &= \langle\phi_-|\phi_+\rangle = \frac{1}{4}\big(\langle B|B\rangle + \langle B|A\rangle - \langle A|B\rangle - \langle A|A\rangle\big) = \frac{i}{2}\,\text{Im}\langle A|B\rangle\,, \\ \rho_{10} &= \rho_{01}^* = -\frac{i}{2}\,\text{Im}\langle A|B\rangle\,.\end{aligned}$$

Hence:

$$\boxed{\rho_{\text{anc}} = \frac{1}{2}\begin{pmatrix}1 + \text{Re}[A(\boldsymbol{\theta})] & i\,\text{Im}[A(\boldsymbol{\theta})] \\ -i\,\text{Im}[A(\boldsymbol{\theta})] & 1 - \text{Re}[A(\boldsymbol{\theta})]\end{pmatrix}}$$

where $A(\boldsymbol{\theta}) := \langle A|B\rangle = \langle\psi_{\text{target}}|U(\boldsymbol{\theta})|\psi_0\rangle$.

Expectation values. *Using $\sigma_Z = \begin{pmatrix}1 & 0 \\ 0 & -1\end{pmatrix}$, $\sigma_Y = \begin{pmatrix}0 & -i \\ i & 0\end{pmatrix}$:*

$$\begin{aligned}\langle\sigma_Z\rangle_{\text{anc}} &= \text{Tr}[\sigma_Z\,\rho_{\text{anc}}] = \rho_{00} - \rho_{11} = \text{Re}[\langle\psi_{\text{target}}|U(\boldsymbol{\theta})|\psi_0\rangle]\,, \\ \langle\sigma_Y\rangle_{\text{anc}} &= \text{Tr}[\sigma_Y\,\rho_{\text{anc}}] = -i\rho_{01} + i\rho_{10} = \text{Im}[\langle\psi_{\text{target}}|U(\boldsymbol{\theta})|\psi_0\rangle]\,.\end{aligned}$$

so $A(\boldsymbol{\theta}) := \langle A|B\rangle = \langle\psi_{\text{target}}|U(\boldsymbol{\theta})|\psi_0\rangle = \langle\sigma_Z\rangle_{anc} + i\langle\sigma_Y\rangle_{anc}$.

***To extract** $\text{Im}[A(\boldsymbol{\theta})]$ **directly from** $\langle\sigma_Z\rangle$**:** insert a phase gate $S = \begin{pmatrix}1 & 0 \\ 0 & i\end{pmatrix}$ on the ancilla after Step 1 (before Step 2). This shifts $|1\rangle \to i\,|1\rangle$, rotating the encoded phase by $\pi/2$ so that $\langle\sigma_Z\rangle = \text{Im}[A(\boldsymbol{\theta})]$.*

Critically, neither $|A\rangle$ nor $|B\rangle$ is measured or collapsed at any point. The complex amplitude $A(\boldsymbol{\theta})$ is encoded coherently in the ancilla Bloch vector, and the work register is returned intact after the circuit.

3.2 Gradient Extraction via Generator Insertion

For each generator G_k, define the augmented COM that extracts the gradient component. Replace $V_A = U(\boldsymbol{\theta})$ with $V_{A_k} = G_k U(\boldsymbol{\theta})$ in Construction 1, giving:

$$A_k(\boldsymbol{\theta}) := \langle\psi_{\text{target}}|G_k U(\boldsymbol{\theta})|\psi_0\rangle\,.$$

Running one base COM (yielding $A(\boldsymbol{\theta})$) and one generator COM (yielding $A_k(\boldsymbol{\theta})$) simultaneously, the CLF gradient signal is:

$$g_k = \frac{\partial F}{\partial \theta_k} = 2\,\mathrm{Im}[A_k(\boldsymbol{\theta})\,A(\boldsymbol{\theta})^*]\,.$$

For a single-Pauli generator $G_k = \frac{1}{2}\sigma_k$:

$$g_k = \mathrm{Im}[\langle \psi_{\text{target}}|\sigma_k U(\boldsymbol{\theta})|\psi_0\rangle \cdot A(\boldsymbol{\theta})^*]\,.$$

Both $A(\boldsymbol{\theta})$ and $A_k(\boldsymbol{\theta})$ are obtained as ancilla expectation values $\langle\sigma_Z\rangle$ and $\langle\sigma_Y\rangle$ from their respective COM instances, without any projective measurement of the work register.

3.3 Coherent Feedback Without Collapse

The CLF architecture does *not* require projective measurement of the ancilla during training. The ancilla expectation values $\langle\sigma_Z\rangle_{\text{anc}}$ and $\langle\sigma_Y\rangle_{\text{anc}}$ are extracted via a pointer-register coupling (Architecture 2 of Section 5.3) or via continuous homodyne detection (Architecture 3), both of which leave the work register intact. The extracted signals feed directly into the classical ODE integrator:

$$\frac{d\theta_k}{dt} = \eta_k\, g_k = 2\eta_k\ \mathrm{Im}[A_k(\boldsymbol{\theta})\,A(\boldsymbol{\theta})^*]\,.$$

3.4 Full COM Circuit Diagram

Base COM — extracts $\mathrm{Re}[A(\boldsymbol{\theta})]$ via $\langle\sigma_Z\rangle$ and $\mathrm{Im}[A(\boldsymbol{\theta})]$ via $\langle\sigma_Y\rangle$:

$$\begin{array}{lcccccccl} |0\rangle_{\text{anc}} & \to & H & \to & \text{ctrl-}V_B^\dagger/\text{ctrl-}V_A & \to & H & \to & \langle\sigma_Z\rangle,\ \langle\sigma_Y\rangle \\ |\psi_0\rangle_w & \to & & \to & V_B^\dagger \text{ or } V_A & \to & & & \text{(work register preserved)} \end{array}$$

Generator COM — same circuit with $V_A \mapsto G_k V_A$:

$$\begin{array}{lcccccccl} |0\rangle_{\text{anc}} & \to & H & \to & \text{ctrl-}V_B^\dagger/\text{ctrl-}(G_k V_A) & \to & H & \to & \langle\sigma_Z\rangle,\ \langle\sigma_Y\rangle \\ |\psi_0\rangle_w & \to & & \to & V_B^\dagger \text{ or } G_k V_A & \to & & & \text{(work register preserved)} \end{array}$$

Im-extraction variant — insert S gate on ancilla after Step 1 to shift $\langle\sigma_Z\rangle \to \mathrm{Im}[A(\boldsymbol{\theta})]$:

$$|0\rangle_{\text{anc}} \ \to\ H \ \to\ S \ \to\ \text{ctrl-ops} \ \to\ H \ \to\ \langle\sigma_Z\rangle = \mathrm{Im}[A(\boldsymbol{\theta})]$$

The two COM instances (base and generator) run in parallel or sequentially to produce $A(\boldsymbol{\theta})$ and $A_k(\boldsymbol{\theta})$, which are fed to the ODE integrator as described in Algorithm 1.

4 Lyapunov Convergence Analysis

4.1 Setup and Lyapunov Function

Define the Lyapunov function:

$$V(\boldsymbol{\theta}) := 1 - F(\boldsymbol{\theta}) = 1 - |\langle\psi_{\text{target}}|U(\boldsymbol{\theta})|\psi_0\rangle|^2 \in [0,1].$$

This satisfies the three Lyapunov conditions:

1. $V(\boldsymbol{\theta}) \geq 0$ for all $\boldsymbol{\theta} \in \mathbb{R}^M$, since $F \in [0, 1]$ by definition of fidelity.
2. $V(\boldsymbol{\theta}) = 0$ if and only if $F(\boldsymbol{\theta}) = 1$, i.e., $U(\boldsymbol{\theta})\,|\psi_0\rangle = e^{i\phi}\,|\psi_{\text{target}}\rangle$.
3. V is smooth (C^∞) as a function of $\boldsymbol{\theta}$ since $U(\boldsymbol{\theta})$ is a smooth unitary and the inner product is smooth.

4.2 The CLF Parameter Update Law

Definition 1 (Coherent Lyapunov Feedback Law). *The CLF update law is the continuous-time ODE:*

$$\frac{d\theta_k}{dt} = \eta_k \frac{\partial F}{\partial \theta_k} = 2\eta_k \, \mathrm{Im}[A_k(\boldsymbol{\theta})\, A(\boldsymbol{\theta})^*]\,,$$

for $k = 1, \dots, M$, where $\eta_k > 0$ are learning rates and $A(\boldsymbol{\theta})$, $A_k(\boldsymbol{\theta})$ are the coherent overlap amplitudes from the COM.

This law is designed so that the time derivative of V is non-positive. Computing $\dot{V}$:

$$\dot{V} = -\dot{F} = -\sum_k \frac{\partial F}{\partial \theta_k} \cdot \frac{d\theta_k}{dt} = -\sum_k \eta_k \left(\frac{\partial F}{\partial \theta_k}\right)^2 \leq 0,$$

with equality iff $\partial F / \partial \theta_k = 0$ for all k, i.e., $\boldsymbol{\theta}$ is a critical point of F.

4.3 Main Convergence Theorem

Theorem 1 (Monotonic Convergence of CLF). *Let $\boldsymbol{\theta}(t)$ evolve under the CLF law (Definition 1) with $\eta_k > 0$. Then:*

(a) $V(\boldsymbol{\theta}(t))$ is monotonically non-increasing: $V(\boldsymbol{\theta}(t')) \leq V(\boldsymbol{\theta}(t))$ for all $t' > t$.

(b) The trajectory $\boldsymbol{\theta}(t)$ converges to the largest invariant set

$$\mathcal{S}^* := \left\{ \boldsymbol{\theta} : \frac{\partial F}{\partial \theta_k} = 0 \text{ for all } k \right\}.$$

(c) If $F = 1$ is the unique maximum of F on the manifold $\mathcal{U}(M)$ (which holds generically for expressive PQCs), then $\boldsymbol{\theta}(t) \to \boldsymbol{\theta}^*$ with $F(\boldsymbol{\theta}^*) = 1$.

Proof sketch: *Part (a) follows from $\dot{V} = -\sum_k \eta_k \left(\partial F / \partial \theta_k\right)^2 \leq 0$. Part (b) follows from LaSalle's invariance principle applied to the compact sublevel sets of V. Part (c) requires genericity of $F = 1$ as the unique global maximum, discussed in Section 4.4.*

4.4 Fixed Point Topology and Spurious Critical Points

The set $\mathcal{S}^*$ contains both the global optimum ($F = 1$) and potential saddle points or local maxima of F. This is analogous to the barren plateau problem in gradient descent.

However, CLF has a structural advantage: the feedback law is derived from V, a Lyapunov function that is globally defined on the compact parameter manifold. The critical points of F are:

- **Global maximum:** $F = 1$. This is the target; CLF converges here if reachable from initial $\boldsymbol{\theta}$.
- **Saddle points:** $\partial F/\partial\theta_k = 0$ but $F < 1$. CLF may be temporarily trapped but saddle points are generically unstable under perturbation.
- **Local maxima:** $F < 1$ but locally optimal. These are the only genuine traps. For PQCs whose dynamical Lie algebra (DLA) is $\mathfrak{su}(2^n)$ (maximally expressive), it has been shown that F has no local maxima other than the global one.

Corollary 2 (Global Convergence for Maximally Expressive PQCs)**.** *If the dynamical Lie algebra*

$$\mathfrak{g} = \mathrm{span}\{G_k,\ [G_k, G_{k'}],\ \ldots\}$$

equals $\mathfrak{su}(2^n)$, then the fidelity landscape $F(\boldsymbol{\theta})$ has no local maxima other than $F = 1$, and CLF converges to the global optimum from almost all initial conditions $\boldsymbol{\theta}(0)$.

4.5 Convergence Rate

Near the optimum, the CLF dynamics linearize. Setting $\delta_k = \theta_k - \theta_k^*$ and expanding to second order:

$$\frac{d\delta_k}{dt} = -\sum_j \eta_k\, H_{kj}\, \delta_j + \mathcal{O}(|\delta|^2),$$

where $H_{kj} = \partial^2 F/\partial\theta_k\,\partial\theta_j|_{\boldsymbol{\theta}^*}$ is the Hessian of F at the optimum. The convergence rate near the optimum is determined by the smallest eigenvalue $\lambda_{\min}$ of the Hessian matrix H:

$$V(\boldsymbol{\theta}(t)) \ \sim\ V(\boldsymbol{\theta}(0))\,\exp(-2\eta\,\lambda_{\min}\,t) \quad \text{as } t \to \infty.$$

This is exponential convergence near the optimum, which is generically faster than the sub-linear rates achievable by stochastic gradient descent under shot noise.

5 The CLF Algorithm

5.1 Complete Algorithm Specification

5.2 Complexity Analysis

The dominant cost of CLF is in the quantum stage (Steps 5–7). Per iteration we require:

- **1 COM instance** for $A(\boldsymbol{\theta})$: $\mathcal{O}(n)$ qubit operations, 1 ancilla qubit. Circuit depth: $\mathcal{O}(n + \mathrm{depth}(U(\boldsymbol{\theta})))$.
- M **generator-COM instances** for $A_k(\boldsymbol{\theta})$: each adds 1 gate (G_k insertion) to the COM. These can be parallelized if M ancilla qubits are available.

Algorithm 1 Coherent Lyapunov Feedback (CLF)

Require: Initial state $|\psi_0\rangle$, target state $|\psi_{\text{target}}\rangle$, PQC generators $\{G_k\}$, initial parameters $\boldsymbol{\theta}^0 \in \mathbb{R}^M$, learning rates $\{\eta_k\}$, convergence threshold $\varepsilon > 0$, maximum time T.

Precomputation (offline, once):

Compute DLA: $\mathfrak{g} = \text{span}\{G_k,\ [G_k, G_j],\ [[G_k, G_j], G_l],\ \ldots\}$
Verify $\dim(\mathfrak{g}) \leq 4^n$ (finiteness check)
if $\dim(\mathfrak{g}) = 4^n - 1$ (i.e., $\mathfrak{su}(2^n)$) **then**: global convergence guaranteed
else: add perturbation strategy for saddle escape (Section 5.3)
end if
Prepare $|\psi_{\text{target}}\rangle$ in register B (fixed throughout training)

Training loop ($t = 0, 1, 2, \ldots$):

repeat
  [Quantum stage]
  Prepare $|\psi(\boldsymbol{\theta})\rangle = U(\boldsymbol{\theta})\,|\psi_0\rangle$ in register A
  Run COM (Construction 1) $\to$ obtain ancilla encoding $A(\boldsymbol{\theta})$
  for each $k = 1, \ldots, M$ **do**
    Prepare $G_k U(\boldsymbol{\theta})\,|\psi_0\rangle$ in register A' (one extra gate insertion)
    Run generator-COM $\to$ obtain ancilla encoding $A_k(\boldsymbol{\theta})$
  end for
  [Classical stage]
  Extract coherent gradient signals:

$$g_k = 2\,\text{Im}[A_k(\boldsymbol{\theta}) \cdot A(\boldsymbol{\theta})^*] \quad \text{for } k = 1, \ldots, M$$

  Integrate ODE one step (e.g., 4th-order Runge-Kutta):

$$\theta_k(t + dt) = \theta_k(t) + \eta_k \cdot g_k \cdot dt$$

  [Optional convergence check]: measure $F(\boldsymbol{\theta})$ via single SWAP test ($\mathcal{O}(1)$ shots)
until $1 - F(\boldsymbol{\theta}) < \varepsilon$ **or** $t > T$

Ensure: Optimal parameters $\boldsymbol{\theta}^*$ with $U(\boldsymbol{\theta}^*)\,|\psi_0\rangle \approx |\psi_{\text{target}}\rangle$

- $\mathcal{O}(1)$ **SWAP test** for convergence check (optional): 1 ancilla, $\mathcal{O}(n)$ depth, $\mathcal{O}(\log(1/\varepsilon))$ shots for ε-precision fidelity estimate.

Total qubit overhead: 1 ancilla (serial) or M ancillae (parallel) plus $2n$ qubits (n for A, n for B). Total projective measurements per convergence: $\mathcal{O}(\log(1/\varepsilon))$ optional shots (vs. $\mathcal{O}(M \times S)$ required shots for parameter shift).

Method	**Shots/step**	**Projective?**	**Barren plateau?**	**Global conv.?**
Parameter Shift	$\mathcal{O}(M \times S)$	Yes, required	Yes	No
SPSA	$\mathcal{O}(S)$	Yes, required	Yes (reduced)	No
Rotosolve	$\mathcal{O}(M \times S)$	Yes, required	No	Local only
CLF (this work)	$\mathcal{O}(S)$	No	No	Yes (expressive)

5.3 Saddle Escape Strategy

When the DLA is not maximally expressive, CLF may stall at saddle points. We propose two mitigation strategies:

- **Random perturbation:** if $\|\mathbf{g}\|^2 = \sum_k g_k^2 < \delta$ for $\delta > 0$, inject a random displacement $\boldsymbol{\theta} \to \boldsymbol{\theta} + \boldsymbol{\xi}$ where $\boldsymbol{\xi}$ is drawn uniformly on a small sphere of radius r. This generically escapes saddle points with probability 1.
- **Momentum augmentation:** add a momentum term m_k to the ODE:

$$\frac{d\theta_k}{dt} = \eta_k g_k + \gamma m_k, \qquad \frac{dm_k}{dt} = g_k - \beta m_k.$$

  This provides inertia that carries trajectories through saddle regions, analogous to heavy-ball methods in classical optimization.

6 Noise Robustness Analysis

6.1 Lindblad Dynamics

On real quantum hardware, the circuit state evolves under the Lindblad master equation rather than pure unitary dynamics:

$$\frac{d\rho}{dt} = -i\big[H_{\text{eff}}(\boldsymbol{\theta}),\, \rho\big] + \sum_j \gamma_j \left(L_j \rho L_j^\dagger - \frac{1}{2}\left\{ L_j^\dagger L_j,\, \rho \right\} \right),$$

where $\{L_j\}$ are jump operators encoding decoherence channels (amplitude damping, dephasing, etc.) and $\gamma_j > 0$ are decay rates. Under Lindblad dynamics the fidelity $F(\boldsymbol{\theta})$ becomes a mixed-state fidelity $F_{\text{mixed}}(\boldsymbol{\theta}) = \langle\psi_{\text{target}}|\, \rho(\boldsymbol{\theta})\, |\psi_{\text{target}}\rangle$.

6.2 Lyapunov Function Under Noise

Define the noisy Lyapunov function:

$$V_{\text{noise}}(\boldsymbol{\theta}) := 1 - \mathrm{Tr}[|\psi_{\text{target}}\rangle \langle\psi_{\text{target}}| \, \rho(\boldsymbol{\theta})] \, .$$

The time derivative under Lindblad dynamics is:

$$\dot{V}_{\text{noise}} = -\sum_k \eta_k \left(\frac{\partial F_{\text{mixed}}}{\partial \theta_k}\right)^2 + D(\boldsymbol{\theta}),$$

where $D(\boldsymbol{\theta})$ is the decoherence-induced drift:

$$D(\boldsymbol{\theta}) = \sum_j \gamma_j \, \mathrm{Tr}\left[|\psi_{\text{target}}\rangle \langle\psi_{\text{target}}| \left(L_j \rho L_j^\dagger - \frac{1}{2}\left\{L_j^\dagger L_j, \, \rho\right\}\right)\right].$$

Theorem 3 (Noise-Robust Convergence)**.** *CLF converges to the vicinity of the optimum under Lindblad noise if the feedback rate dominates decoherence:*

$$\sum_k \eta_k \left(\frac{\partial F}{\partial \theta_k}\right)^2 \geq D(\boldsymbol{\theta}) + \delta \quad \textit{for some } \delta > 0.$$

The steady-state fidelity satisfies:

$$F_{\text{ss}} \geq 1 - \frac{\max_{\boldsymbol{\theta}} D(\boldsymbol{\theta})}{\delta + \sum_k \eta_k \lambda_{\min}^2}.$$

In the limit $\gamma_j \to 0$, $F_{\text{ss}} \to 1$, recovering noiseless convergence.

This theorem quantifies the fundamental trade-off: faster feedback (larger η_k) and stronger overlap signal (larger $\partial F / \partial \theta_k$) fight against decoherence. For systems with coherence times $T_2 \gg 1/\eta$ (i.e., the feedback is fast relative to decoherence), near-unit fidelity is achievable.

6.3 Shot Noise in the COM

While CLF eliminates projective measurement from the optimization loop, the COM interface introduces a classical readout step when extracting the coherent signal g_k. If this extraction involves any measurement (e.g., weak measurement of ancilla), shot noise contributes variance:

$$\mathrm{Var}[g_k] = \frac{1 - g_k^2}{S_{\text{anc}}},$$

where S_{anc} is the number of ancilla readings. However, S_{anc} is a fixed overhead (1 ancilla qubit, $\mathcal{O}(1)$ reads per step) independent of M, giving a total shot cost of $\mathcal{O}(1)$ per iteration — a factor of $\mathcal{O}(M)$ improvement over parameter shift.

7 Extensions and Open Problems

7.1 Lie-Algebraic Circuit Design for CLF

CLF is most efficient when the generators $\{G_k\}$ have a known, small DLA. Circuits with DLA $= \mathfrak{su}(2^n)$ are maximally expressive (guaranteeing global convergence) but require exponentially many parameters. An optimal design target is a DLA that is:

- Large enough to contain the target $U^* = U(\boldsymbol{\theta}^*)$ in its group manifold (expressivity).
- Small enough that the Fourier spectrum of $F(\boldsymbol{\theta})$ is sparse, enabling joint multi-parameter fitting with compressed sensing ($\mathcal{O}(\text{sparsity} \times \log M)$ measurements).

7.2 Compressed Sensing on the Loss Landscape

Because $F(\boldsymbol{\theta})$ is a trigonometric polynomial with frequency content determined by the DLA spectrum, and because the CLF gradient $g_k = \partial F/\partial\theta_k$ inherits this Fourier structure, one can in principle fit the entire gradient vector field with $\mathcal{O}(K \log M)$ COM evaluations, where K is the sparsity of the Fourier representation. This would reduce the quantum stage cost from $\mathcal{O}(M)$ (serial generator-COMs) to $\mathcal{O}(K \log M)$, a significant improvement for sparse circuits.

7.3 CLF for Mixed State Targets

When the target is a mixed state ρ_{target} (e.g., a thermal state or a partially entangled state), the fidelity generalizes to $F = \text{Tr}[\rho_{\text{target}}\, \rho(\boldsymbol{\theta})]$. The CLF law extends naturally with the replacement $|\psi_{\text{target}}\rangle \langle\psi_{\text{target}}| \to \rho_{\text{target}}$ in all overlap computations. The Lyapunov proof carries through with $V = 1 - \text{Tr}[\rho_{\text{target}}\, \rho(\boldsymbol{\theta})]$, which remains non-negative by the quantum data processing inequality.

7.4 Connection to Quantum Optimal Control

CLF is a special case of quantum optimal control (QOC) with the Lyapunov cost function. The connection to GRAPE (Gradient Ascent Pulse Engineering) is particularly close: GRAPE uses parameter-shift-like finite differences on control pulses, while CLF replaces this with a coherent feedback signal. Hybrid CLF-GRAPE algorithms, where CLF drives coarse convergence and GRAPE refines near the optimum, represent a promising direction.

8 Discussion

CLF represents a qualitative departure from measurement-based QML training. The key conceptual shift is from "evaluate loss, compute gradient, take step" to "maintain coherent feedback, flow continuously toward target." This shift has several important implications:

On barren plateaus. The barren plateau phenomenon arises because the gradient of the loss, when computed by parameter shift, has variance that decreases exponentially with system size. CLF avoids this because the feedback signal is not a gradient estimate subject to shot noise — it is a coherent amplitude that encodes the exact discrepancy between current and target states. The signal-to-noise ratio of the COM does not depend on the number of parameters M.

On hardware requirements. The COM requires an additional n-qubit register to hold $|\psi_{\text{target}}\rangle$ and 1–2 ancilla qubits. This doubles the qubit count but removes the measurement overhead entirely. For near-term hardware, this is a favorable trade-off when M is large, since qubit count is increasing faster than coherence time.

On the role of the target state. CLF requires $|\psi_{\text{target}}\rangle$ to be explicitly preparable on the quantum computer. This is the case for quantum compilation, state preparation from classical description, and quantum chemistry ground state problems. For problems where the target is implicitly defined (classification), CLF must be adapted to use a set of labeled target states in a supervised feedback loop.

9 Conclusion

We have presented Coherent Lyapunov Feedback (CLF), a measurement-free training algorithm for parameterized quantum circuits. The key contributions are:

1. The **Coherent Overlap Module**, a quantum circuit that computes fidelity amplitudes without wavefunction collapse.
2. The **CLF update law**, derived from Lyapunov's second method, with guaranteed monotonic convergence.
3. A **rigorous convergence theorem** for expressive PQCs.
4. **Noise robustness analysis** under Lindblad dynamics.
5. A **complete algorithm** with $\mathcal{O}(S)$ measurement overhead, reducing the measurement bottleneck of parameter shift by a factor of $\mathcal{O}(M \times S)$.

CLF opens a new direction for quantum machine learning that draws on the rich theory of quantum feedback control. We believe the intersection of Lyapunov control theory, Lie-algebraic circuit design, and coherent quantum feedback is a fertile ground for future algorithmic development, and we hope this paper serves as a foundation for that work.

References

[1] M. Schuld, V. Bergholm, C. Gogolin, J. Izaac, N. Killoran. Evaluating analytic gradients on quantum hardware. *Physical Review A* 99, 032331 (2019).

[2] D. Wierichs, J. Izaac, C. Wang, C. Y.-Y. Lin. General parameter-shift rules for quantum gradients. *Quantum* 6, 677 (2022).

[3] A. Abbas, R. King, H.-Y. Huang, W. J. Huggins, R. Movassagh, D. Gilboa, J. McClean. On quantum backpropagation, information reuse, and cheating measurement collapse. *NeurIPS* 36 (2024).

[4] K. Chinzei, S. Yamano, Q. H. Tran, Y. Endo, H. Oshima. Trade-off between gradient measurement efficiency and expressivity in deep quantum neural networks. *npj Quantum Information* 11, 79 (2025).

[5] J. R. McClean, S. Boixo, V. N. Smelyanskiy, R. Babbush, H. Neven. Barren plateaus in quantum neural network training landscapes. *Nature Communications* 9, 4812 (2018).

[6] H. M. Wiseman, G. J. Milburn. *Quantum Measurement and Control*. Cambridge University Press (2010).

[7] S. Lloyd. Coherent quantum feedback. *Physical Review A* 62, 022108 (2000).

[8] M. Mirrahimi, R. Van Handel. Stabilizing feedback controls for quantum systems. *SIAM Journal on Control and Optimization* 46, 445–467 (2007).

[9] D. Dong, I. R. Petersen. Quantum control theory and applications: A survey. *IET Control Theory & Applications* 4, 2651–2671 (2010).

[10] M. Larocca, P. Czarnik, K. Sharma, G. Muraleedharan, P. J. Coles, M. Cerezo. Diagnosing barren plateaus with tools from quantum optimal control. *Quantum* 6, 824 (2022).

[11] S. Khatri, R. LaRose, A. Poremba, L. Cincio, A. T. Sornborger, P. J. Coles. Quantum-assisted quantum compiling. *Quantum* 3, 140 (2019).

[12] N. Khaneja, T. Reiss, C. Kehlet, T. Schulte-Herbruggen, S. J. Glaser. Optimal control of coupled spin dynamics. *Journal of Magnetic Resonance* 172, 296–305 (2005).

[13] M. Cerezo, A. Arrasmith, R. Babbush, S. C. Benjamin, S. Endo, K. Fujii, et al. Variational quantum algorithms. *Nature Reviews Physics* 3, 625–644 (2021).

[14] H.-Y. Huang, R. Kueng, J. Preskill. Predicting many properties of a quantum system from very few measurements. *Nature Physics* 16, 1050–1057 (2020).

[15] A. Skolik, J. R. McClean, M. Mohseni, P. van der Smagt, M. Leib. Layerwise learning for quantum neural networks. *Quantum Machine Intelligence* 3, 5 (2021).