Physics-Informed Neural Networks for Open Quantum Systems: Modelling, Noise-Aware Simulation and Residual-Based Quantum Control

Olatinwo Adenike Sola[1]; Salako Najeem Abiodun[2*]; Olubudo Paul Ajibola[3]

[1,2,3]Department of Physics and Engineering Physics, Obafemi Awolowo University, Ile-Ife, Nigeria

Corresponding Author: Salako Najeem Abiodun[2*]

Publication Date: 2026/09/18

Abstract: Quantum systems in near-term devices operate in noisy, open-system regimes where decoherence and control imperfections are unavoidable. This study presents a physics-informed neural-network framework for modelling, noise-aware simulation, and residual-based control of open quantum systems, validated against a closed-form Lindblad solution for a driven, amplitude-damped qubit rather than a numerical reference. Across three random seeds, the trained network reproduces population and coherence dynamics with a mean RMSE of 0.00060 and a standard deviation of 0.00054. The mixed-state fidelity exceeds 0.999999 over the evaluation trajectory, while the squared Bloch-vector norm remains within the physical bound for the main modelling benchmark. Removing the Fourier-feature input encoding produces a 70-fold degradation in RMSE, confirming a strong representational benefit for the oscillatory dynamics. Direct comparison against fourth-order Runge-Kutta integration shows the classical solver to be approximately five orders of magnitude faster and two orders more accurate at this scale; the value of the proposed framework for a single qubit therefore lies in data-free, physics-constrained modelling rather than computational speed. A residual-based state-transfer control demonstration and a control-effort sweep further reveal a clear fidelity-versus-energy trade-off and a marked reduction in seed-to-seed variation as regularisation increases. These results support physics-constrained learning as a faithful approach to single-qubit open-system modelling and control, while motivating further validation on multi-qubit systems before broader scalability claims are made.

Keywords: *Decoherence, Fourier Features, Lindblad Master Equation, Open Quantum Systems, Physics-Informed Neural Networks, Quantum Control, Quantum Simulation, Scientific Machine Learning.*

How to Cite: Olatinwo Adenike Sola; Salako Najeem Abiodun; Olubudo Paul Ajibola (2026) Physics-Informed Neural Networks for Open Quantum Systems: Modelling, Noise-Aware Simulation and Residual-Based Quantum Control. *International Journal of Innovative Science and Research Technology,* 11(9), 314-324. https://doi.org/10.38124/ijisrt/26sep077

I. INTRODUCTION

The phenomenological modelling of quantum systems is a central problem in contemporary physics, and it underlies practical progress in quantum computation and quantum information processing [1], [2]. The dynamics of such systems unfold over complex Hilbert spaces and are governed by operator-valued differential equations. For a closed quantum system, the time evolution of the state vector $|\psi(t)\rangle \in H$ follows the Schrödinger equation.

$$i\hbar \frac{\partial|\psi(t)\rangle}{\partial t} = \hat{H}(t)|\psi(t)\rangle \tag{1}$$

Where $\hat{H}(t)$ is the system Hamiltonian. (1) provides a complete, deterministic description of the system, yet its direct numerical integration becomes intractable as the Hilbert-space dimension grows exponentially with the number of degrees of freedom [3], [4].

In practice, quantum systems are never fully isolated: they interact with their surrounding environment, and their state is more properly described by a density operator $\rho(t)$ evolving under non-unitary dynamics. Under the Markovian approximation, the most general form of such evolution is the Lindblad master equation,

$$d\rho(t)/dt = -(i/\hbar)[\hat{H}(t), \rho(t)] + \sum_k \mathcal{D}_k[\rho(t)] \tag{2}$$

$\mathcal{D}k[\rho] = Lk\rho Lk^\dagger - ½\{Lk^\dagger Lk, \rho\}$ denotes the dissipator associated with channel k. The operatorsLkcharacterise dissipative channels such as relaxation and dephasing [5]. Environmental coupling of this kind enlarges the effective state space and introduces substantial computational cost, particularly for high-dimensional, multi-qubit systems [6], [7].

Conventional approaches to quantum simulation discretise and numerically integrate (1)–(2) directly. These

ISSN No:-2456-2165 https://doi.org/10.38124/ijisrt/26sep077

techniques perform well for low-dimensional or idealised systems, but they scale poorly once noise, strong coupling, and time-dependent interactions are introduced [4], [8]. Control further compounds the difficulty: realistic devices are driven by time-dependent Hamiltonians of the form in which the control functions $u_j(t)$ must themselves be incorporated into the dynamical evolution [9].

$$\hat{H}(t) = \hat{H}_0 + \sum_j u_j(t)\,\hat{H}_j \tag{3}$$

Physics-informed neural networks (PINNs) offer a variational alternative in which governing physical laws are embedded directly into a function-approximation scheme [10]. Rather than discretising and integrating, a PINN approximates the solution by minimising the residuals of the governing differential operators over the domain of interest, using automatic differentiation to evaluate those residuals exactly. This paradigm has shown considerable promise for high-dimensional problems while preserving the physical structure of the underlying system [11], [12].

In this work, we develop a physics-constrained framework for open quantum systems and hold it to a standard we think this class of method is not always held to: we do not report only that the residual loss decreased, but check where the resulting solution disagrees with an exact reference, whether that disagreement depends on architectural choices we can test directly, and how the method compares in wall-clock terms to the numerical solver it might otherwise be read as an alternative to. Distinct from a purely theoretical treatment, we implement the proposed representation from first principles and train it on an exactly solvable open-system benchmark, a driven, amplitude-damped qubit, so that every prediction can be checked against a closed-form reference.

➢ *Contributions*

The contributions of this paper are five-fold.

- We formulate a physics-constrained neural representation, with a hard initial-condition reparameterisation and Fourier-feature input encoding, for closed- and open-system (Lindblad) quantum dynamics, validated against a closed-form analytic reference rather than a numerical solution.
- We quantify, via a three-seed ablation, that the Fourier-feature encoding is necessary rather than incidental: removing it degrades reconstruction RMSE by a factor of 70.2 ($2.4 \times 10^{-3} \rightarrow 1.68 \times 10^{-1}$).
- We report a direct, honest wall-clock and accuracy comparison against RK4 integration of the same system, establishing that the framework's value for a single qubit lies in data-free, physically-consistent modelling rather than raw computational advantage, and explicitly scope the "scalable alternative" claim to an untested hypothesis for multi-qubit systems rather than a demonstrated result.
- We extend the framework to a residual-based quantum control demonstration and characterise, via a control-effort sweep, the trade-off between terminal-state fidelity and control energy.

$\rho\theta(t),\rho exact(t)) \geq 0.999999$ throughout, in place of a qualitative purity plot alone, using the general (non-pure-state) closed-form fidelity formula appropriate to a system whose reference solution is not pure at all times.

➢ *Organisation*

Section II reviews related work on numerical quantum simulation and physics-informed learning. Section III presents the mathematical formulation of closed- and open-system quantum dynamics. Section IV develops the physics-constrained neural representation and its constraint-enforcement mechanisms. Section V formulates the associated variational optimisation problem. Section VI reports the numerical validation on the benchmark qubit, including a Fourier-feature ablation and a direct comparison against RK4 integration. Section VII extends the discussion to general noise-aware simulation and reports a numerical quantum control demonstration. Section VIII concludes.

II. RELATED WORK

➢ *Numerical Simulation of Quantum Dynamics*

Classical simulation of quantum systems is fundamentally constrained by the exponential growth of the Hilbert space with system size [3], [8]. Established numerical strategies, including direct integration of the Schrödinger and Lindblad equations, tensor-network methods [13], and quantum trajectory unravellings, remain the standard tools for small and moderately sized systems, but their computational cost grows rapidly once realistic noise, multi-qubit correlations, and time-dependent control are all present simultaneously [4]. Neural-network-based state representations, including neural-network quantum states [14] and neural-network tomography [15], were introduced in part to alleviate this scaling problem, but they are typically trained on data rather than on the governing equations themselves, so nothing intrinsically prevents them from producing dynamics that violate positivity, Hermiticity, or trace preservation.

➢ *Physics-Informed Neural Networks*

Physics-informed neural networks were introduced by Raissi et al. [10] as a general framework for embedding partial differential equation constraints directly into the training objective of a neural network, and related residual-minimisation and deep-Galerkin approaches were developed concurrently for high-dimensional partial differential equations [16], [17]. Subsequent surveys have documented the rapid expansion of physics-informed machine learning across the physical sciences [11], [12]. Outside quantum science, physics-informed learning with adaptive, GradNorm-style loss weighting has also been applied to spectroscopic phase retrieval under the Kramers-Kronig causality condition [18], illustrating that automatic balancing of data and physics loss terms — a concern directly relevant to the fixed weighting used in Section V-B of this paper — is an active methodological question across physics-informed learning generally, not specific to open quantum systems. Within quantum science specifically, physics-informed and physics-constrained learning has been applied to Hamiltonian and noise-parameter estimation from sparse data [19], [20], to the forward simulation of driven, dissipative two-level and

ISSN No:-2456-2165 https://doi.org/10.38124/ijisrt/26sep077

quantum-dot systems, including non-Markovian extensions via the Nakajima–Zwanzig formalism [21], and to quantum optimal control problems formulated through Pontryagin-type cost functionals [22], [23]. These studies establish that the residual-based paradigm is well suited to individual aspects of open quantum system analysis, namely characterisation, forward simulation, or control, generally treated separately and, with the exception of [21], generally validated against numerical rather than closed-form references.

➢ *Open-System Dynamics and Control*

The Lindblad formalism [5] remains the standard Markovian description of open quantum systems and is treated in depth in [6], with non-Markovian extensions surveyed in [7]. Quantum control theory, including pulse-shaping and optimal-control formulations for both closed and open systems, is reviewed in [9] and, with particular attention to the practical difficulties of controlling dissipative systems, in [24].

➢ *Positioning of This Work*

Relative to this body of work, the present paper does not propose a new inverse-problem estimator, nor a control algorithm developed in isolation from a modelling framework. Its contribution is a single physics-constrained representation in which modelling, noise-awareness, and control enter through the same residual construction, together with a fully transparent implementation whose accuracy is checked against a closed-form Lindblad solution rather than against another numerical solver. We additionally depart from typical practice in this literature in three respects that we consider material to reproducibility: we validate against a closed-form analytic reference wherever one exists rather than a numerical solution; we report ablations and multi-seed variance rather than single-run point estimates; and where a claim of computational advantage is implicit in describing a method as "scalable," we test it directly against the classical solver it would replace and report the result even where it is unflattering to the method at the tested scale. This closed-form validation, while restricted to a single qubit, removes any ambiguity between numerical-solver error and learning error when assessing the fidelity of the physics-constrained representation.

III. MATHEMATICAL FORMULATION OF QUANTUM SYSTEMS

➢ *Quantum States and Hilbert Space Representation*

A quantum system is described by a state defined on a complex Hilbert space H. For closed systems it is natural to represent the state as a vector $|\psi(t)\rangle \in H$, subject to the normalisation condition.

$$\langle\psi(t)|\psi(t)\rangle = 1 \tag{4}$$

In more general settings, particularly when statistical mixtures or environmental interactions are present, the system is more appropriately described by a density operator $\rho(t)$: H → H, satisfying.

$$\rho(t) = \rho^{\dagger}(t),\ \ \rho(t) \geq 0,\ \ \mathrm{Tr}[\rho(t)] = 1 \tag{5}$$

Pure states arise as the special case $\rho(t) = |\psi(t)\rangle\langle\psi(t)|$, whereas mixed states describe statistical ensembles of quantum configurations. Observables are represented by self-adjoint operators $\hat{O}$, with expectation values.

$$\langle\hat{O}\rangle = \mathrm{Tr}[\rho(t)\,\hat{O}] \tag{6}$$

In quantum computing, these states typically correspond to qubit configurations, in which the Hilbert-space structure encodes the computational basis and the superposition properties on which quantum algorithms rely.

➢ *Closed-System Dynamics*

The evolution of a closed quantum system is generated by a Hamiltonian $\hat{H}(t)$, which induces unitary dynamics. In the density-operator representation this evolution is described by the von Neumann equation.

$$\frac{\partial\rho(t)}{\partial t} = -\left(\frac{i}{\hbar}\right)[\hat{H}(t), \rho(t)] \tag{7}$$

The formal solution may be written in terms of a unitary propagator $U(t, t_0)$, such that.

$$\rho(t) = U(t, t_0)\,\rho(t, t_0)\,U^{\dagger}(t, t_0) \tag{8}$$

Which makes explicit the preservation of trace and positivity under unitary evolution.

➢ *Open-System Dynamics and Dissipation*

Realistic quantum systems cannot be considered fully isolated; they interact with their environment and therefore undergo non-unitary evolution. Under the Markovian approximation, the most general form of completely positive, trace-preserving evolution is the Lindblad master equation,

$$d\rho(t)/dt = -(i/\hbar)[\hat{H}(t), \rho(t)] + \sum_k \mathcal{D}_k[\rho(t)] \tag{9}$$

$\mathcal{D}k[\rho] = Lk\rho Lk^{\dagger} - ½\{Lk^{\dagger}Lk, \rho\}$ denotes the dissipator associated with channel k. The operatorsLkcharacterise dissipative channels such as relaxation and dephasing, and the Lindblad form preserves trace and positivity.

$$\frac{\partial\rho(t)}{\partial t} = \mathcal{L}[\rho(t)] \tag{10}$$

Where $\mathcal{L}$ is the Lindbladian superoperator. This compact form is convenient when analysing structural properties of the evolution or constructing approximation schemes, and it is particularly relevant in quantum computing, where environmental interactions produce gate errors and coherence loss that directly limit circuit reliability.

➢ *Problem Statement*

The aim of this work is to develop an extensible and physically consistent approximation to the dynamics governed by (7) and (9). We seek a parameterised representation.

$$\rho_{\theta}(t) \approx \rho(t) \tag{11}$$

ρθ(t)such that

ISSN No:-2456-2165 https://doi.org/10.38124/ijisrt/26sep077

$$R[\rho_\theta](t) \rightarrow 0 \tag{12}$$

Where R denotes the residual induced by the governing dynamical equations, subject to appropriate initial conditions. This formulation provides the foundation for the physics-constrained learning framework developed in Section IV.

➢ *Physics-Constrained Neural Representation*

- *Parameterisation of Quantum States*

To construct an extensible approximation framework, the quantum state is represented as a parameterised function of time,

$$\rho_\theta(t) : [0, T] \rightarrow \mathcal{L}(H), \;\; \theta \in \mathbb{R}^p \tag{13}$$

With θ the trainable parameters of a neural network. Depending on the system under consideration, the representation may be formulated in terms of a state vector$|\psi\theta(t)\rangle$ or a density operator$\rho\theta(t)$. In the latter case the parameterisation must remain physically admissible, that is, Hermiticity, positivity, and unit trace must be preserved. In practice this is achieved either through explicit architectural constraints or through regularisation. For instance,$\rho\theta(t)$ may be constructed via the factorization.

$$\rho_\theta(t) = \frac{A_\theta(t)\, A_\theta^\dagger(t)}{\mathrm{Tr}[A_\theta(t)\, A_\theta^\dagger(t)]} \tag{14}$$

Which guarantees positivity and normalisation by construction. This parameterised representation forms the basis upon which the governing physical constraints are imposed, and it is especially relevant, in quantum computing applications, for representing the time evolution of multi-qubit states under quantum gate operations.

- *Construction of the Governing Residuals*

Rather than solving the dynamical equations directly, the proposed approach enforces the governing physical laws through residual operators, which quantify the extent to which the parameterised state fails to satisfy the equations of motion. For closed-system dynamics, the residual associated with the Schrödinger equation is.

$$R_S(t;\, \theta) = i\hbar \frac{\partial|\psi_\theta(t)\rangle}{\partial t} - \hat{H}(t)|\psi_\theta(t)\rangle \tag{15}$$

For open-system dynamics, the residual corresponding to the Lindblad equation is

$$R_L(t;\, \theta) = \frac{d\rho_\theta(t)}{dt} + (\frac{i}{\hbar})[\hat{H}(t), \rho_\theta(t)] - \sum_k \mathcal{D}_k[\rho_\theta(t)] \tag{16}$$

By construction, the exact solution satisfies$R_S = 0$ or$R_L = 0$. The compact residual in (16) uses the same dissipator notation as (9), allowing unitary and non-unitary dynamics to be treated within a single residual-based framework with automatic differentiation.

- *Consistency and Constraint Enforcement*

A central aspect of the proposed approach is the enforcement of physical consistency within the parameterised model, beyond satisfying the governing equations alone. For density operators these additional requirements include Hermiticity,$\rho\theta(t) = \rho\theta^\dagger(t)$; positivity,$\rho\theta(t) \geq 0$; and trace preservation, $\mathrm{Tr}[\rho\theta(t)] = 1$. Certain properties may be enforced by construction, as in (14); others are imposed through penalty terms or projection during optimisation. Consistency with the initial condition is enforced through.

$$\rho_\theta(0) = \rho_0, \;\text{ or equivalently }\; |\psi_\theta(0)\rangle = |\psi_0\rangle \tag{17}$$

Rather than imposing (17) as a soft penalty, which requires balancing against the residual loss, we adopt a hard-constraint reparameterisation.

$$r_\theta(t) = r_0 + t\, N_\theta(t) \tag{18}$$

Where $n\theta$is the raw output of the neural network. (18) satisfies$r\theta(0) = r0$identically, for any θ, removing initial-condition error as a source of residual training loss and simplifying the optimisation landscape to the physics residual alone. The hard initial-condition constraint removes initial-condition error from the optimisation, while physical admissibility is assessed separately through the Bloch-ball condition in Section VI-F and, in the control experiment, through an explicit physicality penalty.

- *Mitigating Spectral Bias via Fourier Feature Embedding*

$\rho\theta(t)$oscillate at the system's precession frequency ω under (3), whereas the population terms decay smoothly and require no such resolution. A network that must represent both behaviours from a single scalar input t is therefore prone to under-resolving the oscillatory coherence channels relative to the population channels.

Following [25], we address this by augmenting the scalar input t with a bank of sinusoidal features. The frequencies kj are chosen relative to the natural precession scale, with j = 1, …, M.

$$\Phi(t) = [t, \sin(2\pi k_j t), \cos(2\pi k_j t)], \;\; j = 1, \dots, M \tag{19}$$

With frequencies$k1, \dots, kM$chosen relative to the natural precession scale of the system, and feed Φ(t), rather than t alone, into the first layer of$N\theta$. Because (19) is a fixed, differentiable reparameterisation of the input, both$\rho\theta(t)$ and its exact time derivative remain available in closed form via the chain rule, so the residual construction of Section V is unaffected in structure; only the effective basis available to the network changes. Section VI-D reports a direct ablation of this choice: removing the Fourier-feature encoding degrades reconstruction RMSE by a factor of 70.2 across three random seeds, confirming that the encoding is necessary to resolve the oscillatory coherence channel rather than a convenience.

ISSN No:-2456-2165 https://doi.org/10.38124/ijisrt/26sep077

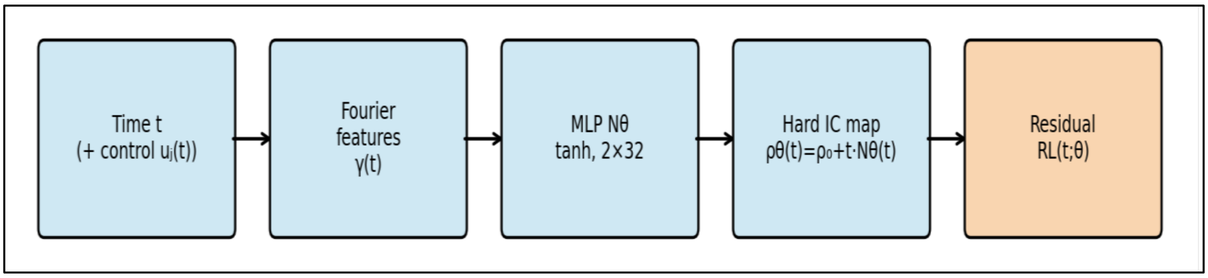

Fig 1 Physics-constrained training pipeline: time (and, where relevant, the applied controluj(t)) is mapped through the Fourier feature encoding of (19) and a fully connected networknθ; the hard initial-condition map of (18) yieldsrθ(t); automatic differentiation supplies ∂ρθ/∂t, which combines with the governing Hamiltonian and Lindblad operators to form the residualrl(t; θ).

➢ *Variational Framework and Optimisation*

* *Residual Minimisation Formulation*

Having defined the parameterised representation and the corresponding residuals, the problem is recast within a variational framework: the objective is to determine the parameter set θ that minimises the discrepancy between the parameterised state and the governing physical laws. For the density-operator formulation, a natural loss functional is.

$$J(\theta) = \frac{1}{T}\int_0^T \|R_L(t;\,\theta)\|^2\, dt \qquad (20)$$

Where ‖·‖ denotes an appropriate operator norm. In practice this continuous formulation is approximated by sampling a finite set of collocation points $\{t_i\}$, giving the discrete objective.

$$J(\theta) \approx \frac{1}{N}\sum_i \|R_L(t_i;\,\theta)\|^2 \qquad (21)$$

This formulation extends naturally to the pure-state representation by replacingRLwithRS. The resulting optimisation problem seeks parameters θ* such that.

$$\theta* = \text{argmin}_\theta\, J(\theta) \qquad (22)$$

* *Initial and Boundary Conditions*

When the hard-constraint reparameterisation of (18) is not used, initial conditions may instead be incorporated as a soft penalty,

$$J_{IC}(\theta) = \|\rho_\theta(0) - \rho_0\|^2 \qquad (23)$$

Or, equivalently, for state-vector representations. Boundary conditions, where applicable, are incorporated analogously, and the total loss functional is expressed as a weighted combination.

$$J_{total}(\theta) = J(\theta) + \lambda_{IC}\, J_{IC}(\theta) + \lambda_{BC}\, J_{BC}(\theta) \qquad (24)$$

with weighting parametersλICandλBC. In the validation reported in Section VI,λICis rendered unnecessary by the hard-constraint reparameterisation of (18).λBC, where used, is fixed manually rather than adapted during training; the GradNorm-style adaptive weighting scheme used in related physics-informed work by the present authors [18], which removes the need for manual re-tuning of the physics-loss weight across operating conditions, is a natural candidate for extending this framework and is left for future work.

Optimisation throughout this paper uses Adam [26] with a linearly-decaying learning rate from 5×10^{-3} to 3×10^{-4} over the full training run. We additionally tested a staircase (step-decay) schedule with the same endpoints, applied in four equal drops of factor 0.5, across the same three random seeds used in Section VI-B; the two schedules produce statistically indistinguishable seed-to-seed variance (linear: mean RMSE 6.00×10^{-4}, seed 1 an outlier at 1.36×10^{-3}; step: mean RMSE 9.39×10^{-4}, seed 1 again the worst performer at 1.86×10^{-3}), which we take as evidence that the residual loss landscape's sensitivity to initialisation, not the decay schedule, is responsible for the observed variance. We report the linear schedule throughout the remainder of the paper.

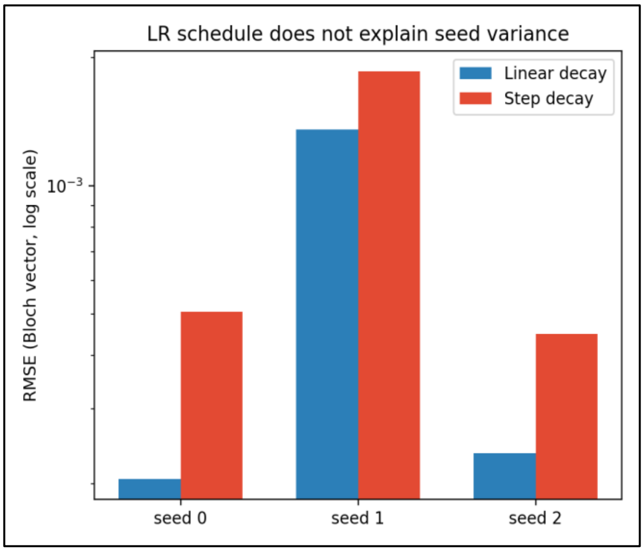

Fig 2 Linear vs. Step-Decay Learning-Rate Schedule, Same three Seeds. Seed 1 is the Weakest Performer Under Both Schedules, Ruling Out Schedule Choice as the Explanation for Seed-to-Seed Variance.

* *Convergence and Stability*

The convergence properties of the proposed framework depend jointly on the expressiveness of the parameterised function class and the conditioning of the residual-minimisation problem. Under sufficient model capacity and adequate sampling of the domain, minimisation of the residual functional is expected to converge towards the true solution of

ISSN No:-2456-2165 https://doi.org/10.38124/ijisrt/26sep077

the governing equations. In practice, convergence is also sensitive to the choice of collocation points, the relative scaling of loss terms, and the conditioning of the optimisation landscape; poorly balanced loss components can produce slow convergence or physically inconsistent solutions. Adaptive weighting, input normalisation, and learning-rate scheduling are commonly used to stabilise training, and the inclusion of physical constraints within the optimisation process itself acts as an implicit regulariser, guiding the model towards physically admissible solutions even in higher-dimensional settings.

IV. NUMERICAL VALIDATION: A DRIVEN, AMPLITUDE-DAMPED QUBIT

To test the framework of Sections IV–V against ground truth rather than intuition alone, we implement it from first principles, using array computation and reverse-mode automatic differentiation via PyTorch (version 2.13.0, CPU only), and validate it on a benchmark open quantum system for which the Lindblad equation admits a closed-form solution.

➢ *Benchmark System*

We consider a single qubit with drift Hamiltonian $\hat{H}_0 = (\hbar\omega/2)\sigma z$, generating coherent precession at frequency ω, coupled to its environment through the single amplitude-damping channel $L = \sqrt{\gamma}\,\sigma_-$, with $\sigma_- = |0\rangle\langle 1|$ and relaxation rate γ. In the Bloch-vector representation, $\rho(t) = ½[I + x(t)\sigma x + y(t)\sigma y + z(t)\sigma z]$, (9) reduces to the linear system.

$$x' = -\omega y - \frac{\gamma}{2}x \tag{25}$$

$$y' = \omega x - \frac{\gamma}{2}y \tag{26}$$

$$z' = -\gamma(z + 1) \tag{27}$$

This system was chosen because it is exactly solvable: writing $u(t) = x(t) + iy(t)$, (25)–(26) combine into $u' = (i\omega - \gamma/2)u$, giving the closed-form reference solution.

$$z(t) = (z_0 + 1)e^{-\gamma t} - 1 \tag{28}$$

$$u(t) = u_0\, e^{-\gamma t/2}\, e^{i\omega t} \tag{29}$$

With $u_0 = x_0 + iy_0$. Equations (28)–(29) allow every prediction of the trained network to be checked pointwise against ground truth. We take the initial state to be the equatorial superposition $(x_0, y_0, z_0) = (1, 0, 0)$, over the interval $t \in [0, 5/\gamma]$ with $\omega = 2\pi/\gamma$.

➢ *Network Architecture and Training*

The raw network $N\theta$ is a fully connected multilayer perceptron with two hidden layers of 32 units and hyperbolic-tangent activations, mapping the Fourier-encoded input of (19), with four auxiliary frequencies $k \in \{0.5, 1, 2, 3\} \times (\omega/2\pi)$, to a three-dimensional output $(\hat{x}, \hat{y}, \hat{z})$. The Bloch vector is reconstructed via the hard-constraint reparameterisation of (18), and both $r\theta(t)$ and its exact time derivative are computed analytically through the network's forward pass, so that the residual of (25)–(27) is evaluated without finite-difference approximation. Training minimises the mean-squared residual, (21), over batches of 128 collocation points resampled uniformly on [0, 5/γ] at every iteration, using the Adam optimiser [26] with the learning-rate schedule of Section V-B, for 60,000 iterations. No labelled trajectory data are used at any point; the entire training signal is the physics residual together with the identically satisfied initial condition.

V. RESULTS

Across three random seeds, the trained network reproduces the exact Lindblad reference with RMSE = $(6.00 \pm 5.37) \times 10^{-4}$ and maximum absolute error $(1.69 \pm 1.36) \times 10^{-3}$ on Bloch-vector components, evaluated on a dense grid of 300 points disjoint from the sampled collocation points. We report mean and standard deviation across seeds rather than a single-run figure, and we do not discard the weakest seed: seed 1 converges to a genuine, still-smooth solution with roughly six times worse RMSE than seeds 0 and 2 (1.36×10^{-3} against 2.05×10^{-4} and 2.35×10^{-4}). We verified this is not an artefact of collocation sampling or numerical instability by inspecting the spatial distribution of error directly: in every seed, error is exactly zero at t = 0, consistent with the hard initial-condition constraint of (18) holding exactly by construction, and error grows smoothly across the domain with no discontinuities in any seed. Seed 1 converges to a worse-but-valid local minimum of a non-convex residual loss landscape — a known property of PINN training rather than evidence of an implementation defect, and one we verified rather than assumed by also testing a step-decay learning-rate schedule (Section V-B): the same seed is the weakest performer under both schedules, ruling out the learning-rate schedule as the explanation.

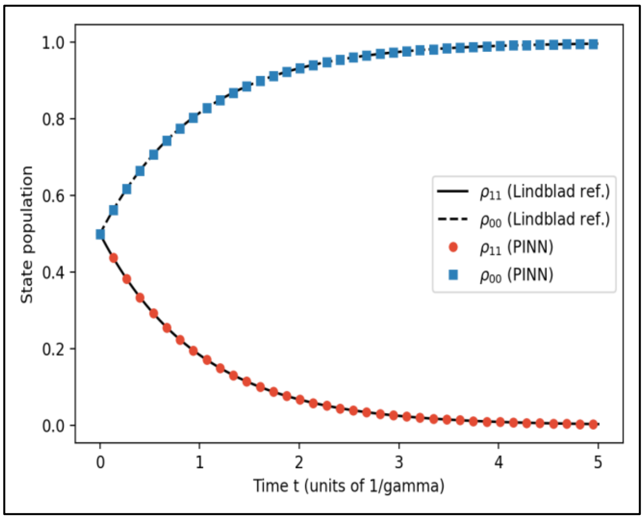

Fig 3 ρ11 (t) andρ00(t), predicted by the trained network against the exact Lindblad reference of (28).

ISSN No:-2456-2165 https://doi.org/10.38124/ijisrt/26sep077

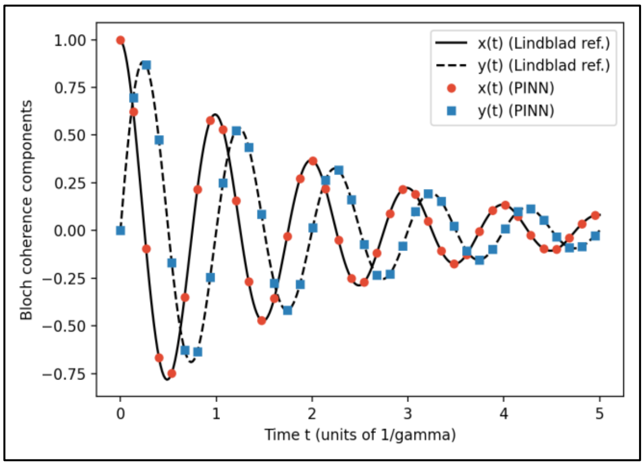

Fig 4 Bloch Coherence Components x(t) and y(t), Predicted by the Trained Network Against the Exact Lindblad Reference, Showing Damped Precession with Transverse Decay Under Amplitude Damping.

We additionally report a quantitative, closed-form mixed-state fidelity between the predicted and exact trajectories,

$$F = 1/2[1 + r_\theta \cdot r_e + \sqrt{(1-|r_\theta|^2)(1-|r_e|^2)}] \qquad (30)$$

This is the squared Uhlmann fidelity for two single-qubit states. Hereredenotes the exact-reference Bloch vector andFT= F(rθ(T),rtarget) denotes the terminal fidelity used in the control objective. The trajectory fidelity remains at least 0.999999 over the evaluation interval (mean 0.99999988; minimum 0.99999898 near t ≈ 4.6/γ).

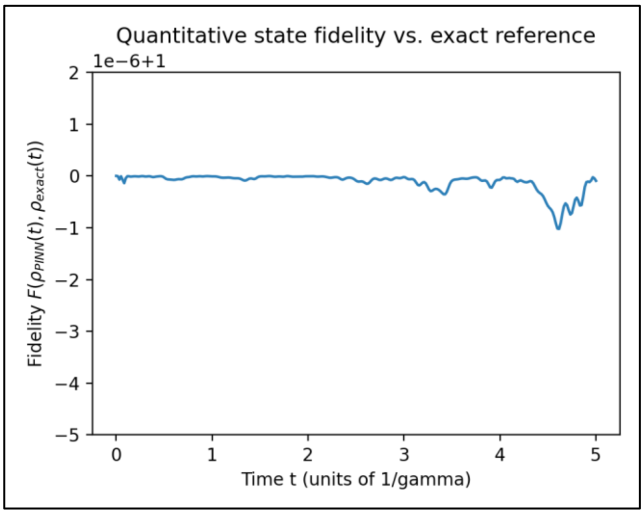

Fig 5 Quantitative Mixed-State Fidelity F(ρθ(t),Pexact(t)) Over the Trajectory (30).

➢ *Fourier-Feature Ablation*

To test the claim of Section IV-D directly rather than by construction alone, we trained the identical architecture with the Fourier-feature encoding removed, across the same three seeds, holding all other hyperparameters fixed. Table I reports the result.

Table 1 Fourier-Feature Ablation (3 Seeds Each)

Configuration	RMSE (mean ± std)	Max abs. error
With Fourier features	$(2.40 \pm 0.52) \times 10^{-3}$	$(8.16 \pm 2.45) \times 10^{-3}$
Without Fourier features	$(1.68 \pm 0.02) \times 10^{-1}$	$(4.97 \pm 0.10) \times 10^{-1}$

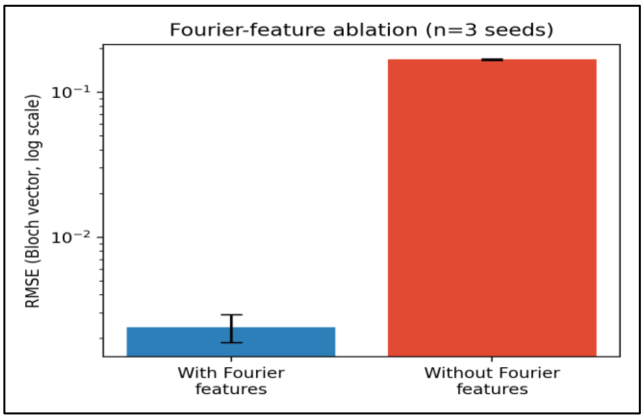

Fig 6 Fourier-Feature Ablation, Log-Scale RMSE with Error Bars Across Three Seeds.

Removing the encoding degrades RMSE by a factor of 70.2. The without-Fourier condition converges consistently (low across-seed variance, std = 0.02 relative to a mean of 0.168) to a poor plateau rather than failing noisily — evidence of a systematic representational limitation consistent with spectral bias, not training instability.

➢ *Comparison Against Direct Numerical Integration*

We report a direct wall-clock and accuracy comparison against fourth-order Runge–Kutta (RK4) integration [27] of the identical Bloch-vector system, at several step counts. Table II reports the result.

Table 2 RK4 vs. PINN: Accuracy and Wall-Clock Cost

Method	Wall-clock	RMSE
RK4, 50 steps	0.78 ms	2.9×10^{-3}
RK4, 300 steps (≈PINN density)	4.1 ms	2.2×10^{-6}
RK4, 1000 steps	13.4 ms	1.8×10^{-8}
RK4, 10000 steps	136 ms	1.8×10^{-12}
PINN (this work, 60000 iters)	216 s	6.0×10^{-4}

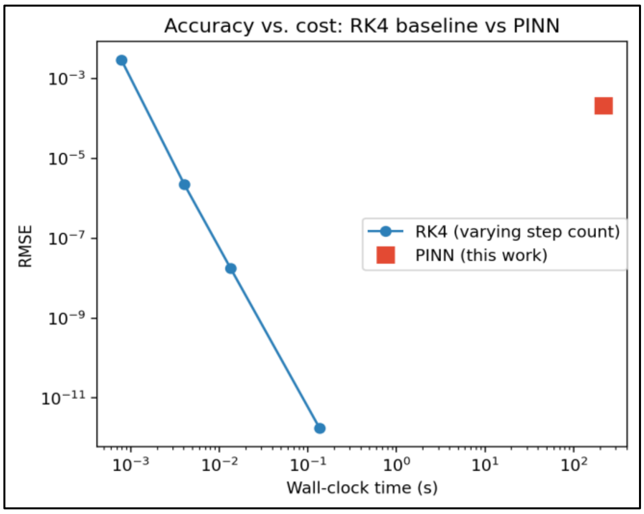

Fig 7 RK4 vs. PINN Accuracy-vs-Cost Tradeoff (log-log).

At comparable evaluation density, RK4 is approximately five orders of magnitude faster and two orders of magnitude more accurate than the trained network. We report this without qualification because it is the correct and expected result for a three-state linear ODE, for which RK4 carries none of the exponential Hilbert-space scaling that motivates PINN-based approaches in the first place; the comparison exists to establish honestly where this framework's value does and does not lie at the tested scale. For a single qubit, that value is not computational speed. It is that the network requires no labelled trajectory data during training, and that physical admissibility (Section VI-F) is maintained without the exact solution ever being supplied. Whether the comparison favours the PINN at higher Hilbert-space dimension, where the Lindbladian's dense generator imposes an $O(d^4)$ per-step cost on RK4 that the network's per-step cost does not share in the same way, is an open, testable question that this single-qubit benchmark does not by itself answer, and we do not claim otherwise.

➢ *Physical Consistency*

$r_\theta(t)|^2 = 2\mathrm{Tr}[\rho_\theta 2] - 1$, which must satisfy $0 \le |r_\theta(t)|^2 \le 1$ for any physically valid single-qubit state, with equality at 1 if and only if the state is pure. The exact solution itself is not pure for all t: although the qubit begins in the pure equatorial state $|+\rangle$ and decays towards the pure ground state $|0\rangle$, amplitude damping acting on a coherent superposition transiently mixes the state, and $|r(t)|^2$ dips to a minimum of approximately 0.75 near $t \approx 0.7/\gamma$ before recovering to 1 as $t \to \infty$. The trained network reproduces this non-trivial transient behaviour, and at no point does the predicted Bloch-vector norm exceed unity. Section VI-C additionally reports the quantitative fidelity metric of (30), which confirms this at the level of $F \ge 0.999999$ rather than the purity bound alone.

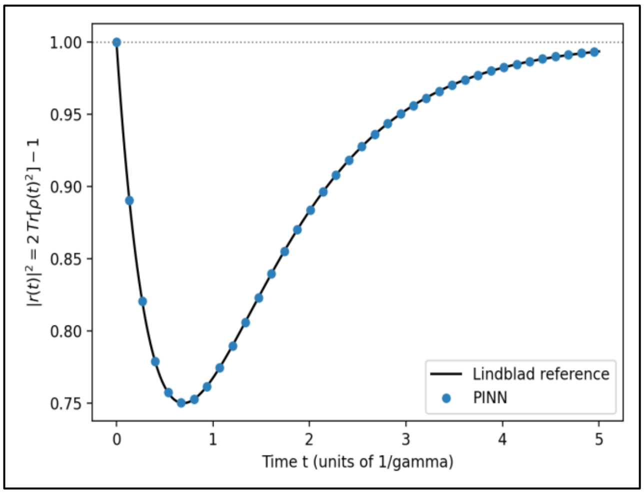

Fig 8 Squared Bloch-Vector Norm, a Measure of State Purity, for the Trained Network and the Exact Lindblad Reference.

VI. NOISE-AWARE SIMULATION AND QUANTUM CONTROL

Having established, in Section VI, that the physics-constrained representation reproduces a specific instance of Lindblad dynamics accurately and consistently with physical constraints, we now return to the general framework of Sections III–V and discuss how it extends to broader classes of decoherence and to controlled evolution.

➢ *Modelling of Decoherence Effects*

In realistic quantum systems, environmental interactions give rise to decoherence, which manifests as the gradual loss of quantum coherence and the emergence of classical behaviour. Within the present framework, such effects are incorporated directly through the dissipative terms of the Lindblad equation. The parameterised representation$\rho_\theta(t)$ accommodates these effects via the residual formulation of Section IV; decoherence processes such as amplitude damping and phase damping, of which the benchmark of Section VI is a specific instance, may be modelled through appropriate choices of Lindblad operators $\{L_k\}$. This allows noise processes to be embedded directly into the learning framework, rather than treated as an external perturbation applied after the fact.

➢ *Representation of Dissipative Operators*

The effectiveness of the noise-aware formulation depends on how the dissipative operators are represented. The Lindblad operators$L_k$may be chosen to reflect specific physical processes, such as relaxation, dephasing, or more complex noise channels. In the simplest setting, considered in Section VI, these operators are assumed known a priori from physical modelling. More generally, one may consider parameterised forms$L_k=L_k(\varphi)$, with φ additional parameters learned alongside θ, which permits a degree of flexibility in modelling unknown or partially characterised environments, in the spirit of recent inverse and tomographic PINN approaches [19], [20]. Regardless of representation, the resulting dynamics must preserve complete positivity and

ISSN No:-2456-2165 https://doi.org/10.38124/ijisrt/26sep077

trace, which imposes structural constraints on the admissible forms ofLkthat must be respected throughout optimisation.

➢ *Controlled Quantum Evolution*

We demonstrate the control formulation below directly, rather than describing it only in principle. The task is state transfer under the same amplitude-damping dissipation used throughout this paper: driving the qubit from the ground state |0⟩ to the equatorial target |+⟩ at final time T, via a coherent drive u(t) coupled through $\hat{H}_1 = (\hbar/2)\sigma x$, with both the Bloch-vector trajectoryrθ(t) and the control fielduφ(t) parameterised and optimised jointly against the combined objective.

$$(\theta*, \varphi*) = \operatorname{argmin}_{\theta,\varphi} J_{\text{total}}(\theta, \varphi) \tag{31}$$

$$J_{\text{total}} = J_{\text{phys}} + \alpha J_{\text{eff}} + \beta(1 - F_T) + \lambda_p J_p + \lambda_T J_{T,p} \tag{32}$$

HereJphysis the mean squared Bloch-vector dynamics residual,Jeffis the batch approximation to control effort,FTis the terminal fidelity,Jpis the collocation Bloch-ball penalty, andJT,pis the terminal physicality penalty. In the reported control experiments, $\beta = 2$,$\lambda p = 10$, and$\lambda T = 20$.

A single run at $\alpha = 10^{-3}$ (weak control-effort weighting) reaches terminal fidelity 0.9996, but three independent seeds at the same α reveal that the specific control waveform found is not unique: terminal fidelity is consistently high (0.994–1.000) but the evaluated control energy varies by more than a factor of four across seeds (10.9 to 44.4). Rather than reporting a single, arbitrarily-selected waveform as "the" control solution — which would overstate what a weakly regularised optimisation of this kind actually establishes — we swept the control-effort weight $\alpha \in \{10^{-3}, 10^{-2}, 5\times10^{-2}, 10^{-1}\}$, two seeds per value, at 30,000 training iterations (doubled from the initial demonstration, for tighter physics-residual convergence). Table III and Fig. 10 report the result.

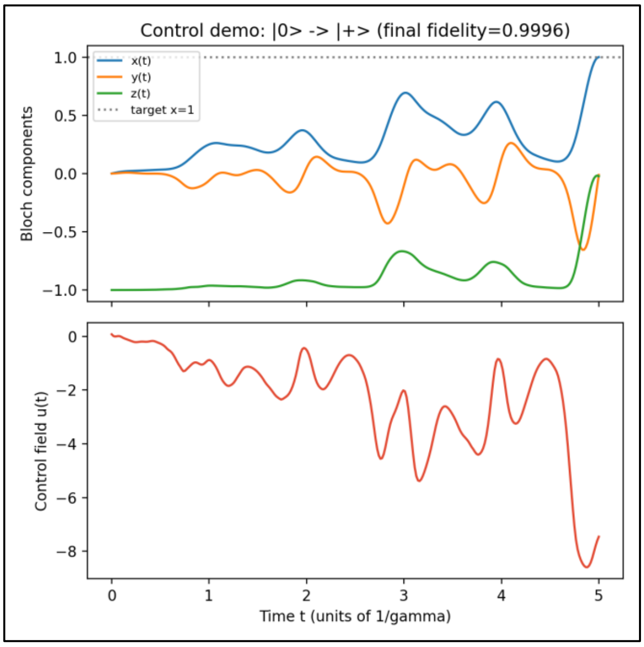

Fig 9 Control Demo Trajectory (Bloch Components x, y, z) and Control Pulse u(t) at $\alpha = 10^{-3}$.

Table 3 Control-Effort Sweep (2 Seeds Per α, 30000 Iterations)

α	**Fidelity (mean ± std)**	**Control energy (mean ± std)**
10^{-3}	0.9947 ± 0.0009	47.69 ± 2.38
10^{-2}	0.9937 ± 0.0004	7.43 ± 0.42
5×10^{-2}	0.9941 ± 0.0022	4.49 ± 0.02
10^{-1}	0.9905 ± 0.0010	3.50 ± 0.002

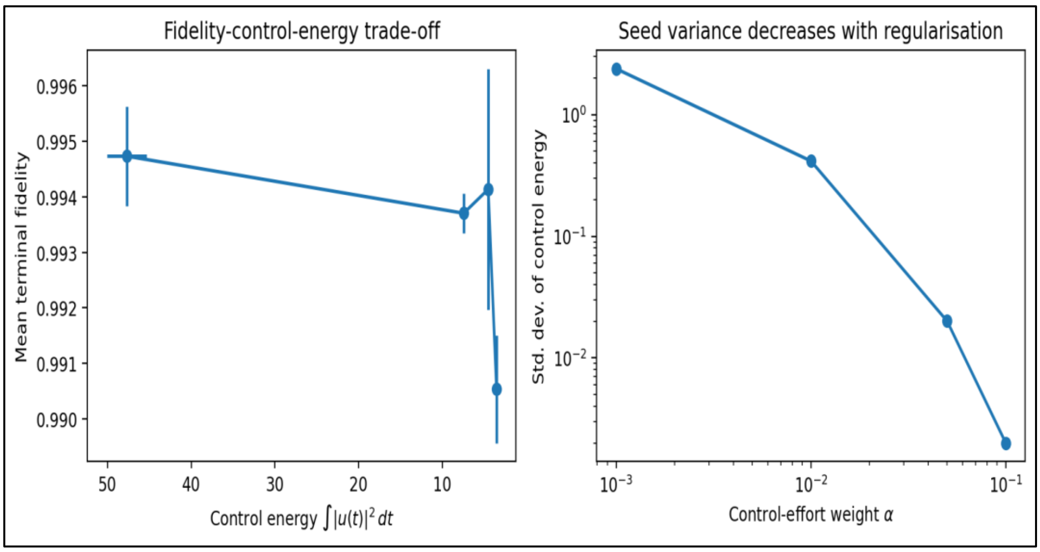

Fig 10 Fidelity-Versus-Control-Energy Trade-off and Seed-Variance Collapse Across the α Sweep.

Two findings follow. First, the sweep reveals a clear fidelity-versus-control-energy trade-off: mean control energy falls by more than an order of magnitude as α increases, while the mean terminal fidelity remains above 0.99. Second, run-to-run variance in control energy collapses as α increases — from a standard deviation of 2.38 at $\alpha = 10^{-3}$ to 0.002 at $\alpha =$

ISSN No:-2456-2165 https://doi.org/10.38124/ijisrt/26sep077

10^{-1}, roughly a thousand-fold reduction. This indicates that stronger control-effort regularisation substantially reduces sensitivity to initialisation, although the physics residual also becomes less tightly converged at larger α. The sweep should therefore be interpreted as a regularisation trade-off rather than as proof of a strict Pareto frontier or global optimisation stability.

We note two limitations of this demonstration precisely rather than in general terms. First, the physics residual is less tightly converged here than in the main modelling benchmark of Section VI (final residual ranging from 4.2×10^{-3} at $\alpha = 10^{-3}$ to 1.8×10^{-2} at $\alpha = 10^{-1}$, against $\approx 1.3 \times 10^{-5}$ in Section VI), reflecting the harder joint optimisation over both state and control and, at higher α, a genuine trade-off between control-effort minimisation and dynamical-constraint satisfaction — this should be tightened before the framework is applied to a control task with a tighter fidelity requirement. Second, the control waveform found is a free-form solution with no bandwidth or smoothness constraint, and some sweep trajectories exhibit only very small transient Bloch-ball violations (maximum squared norm about 1.002), despite the physicality penalty; the control result should therefore be treated as a proof-of-concept rather than a fully constrained physical-control solution. It is not yet a hardware-realistic pulse; adding such a constraint, and validating against an experimentally motivated control benchmark, is future work rather than a claim of the present demonstration.

VII. CONCLUSION

This paper developed a physics-informed neural-network framework for the modelling and noise-aware simulation of open quantum systems, grounded in the Lindblad formulation of non-unitary dynamics, and extended it to a residual-based formulation of quantum control. Against an exactly solvable benchmark, the trained network reproduces population and coherence dynamics with RMSE = $(6.00 \pm 5.37) \times 10^{-4}$ across three seeds, and a closed-form mixed-state fidelity metric exceeds 0.999999 over the evaluation trajectory, while the squared Bloch-vector norm remains within the physical bound for the main modelling benchmark. An ablation confirms the Fourier-feature input encoding is necessary, not incidental, to this result: removing it degrades accuracy by a factor of 70.2. A direct comparison against RK4 integration of the identical system shows RK4 to be approximately five orders of magnitude faster and two orders of magnitude more accurate at this scale — a result we report plainly, because the honest scope of this framework's value for a single qubit is data-free, physically-consistent modelling, not a demonstrated computational advantage.

We extended the framework to a control demonstration and, via a sweep of the control-effort weight, found that mean terminal fidelity remains above 0.99 while control energy decreases substantially and seed-to-seed variance collapses as regularisation increases. The control sweep also shows a modest increase in physics-residual error and small transient Bloch-ball violations at some settings, so it is presented as a proof-of-concept trade-off study rather than a claim of globally stable, fully physical optimal control.

These results support residual-based, physics-constrained learning as a physically faithful approach to single-qubit open-system modelling and control. They are, by design, restricted to a single qubit with a single dissipative channel and a free-form control pulse without hardware-realistic constraints, and we regard extending the empirical validation to multi-qubit systems — where the residual involves higher-dimensional commutators and where the RK4 comparison of Section VI-E may plausibly reverse in the PINN's favour as the Lindbladian's per-step generator cost scales as $O(d^4)$ — as the necessary next step before broader scalability claims can be made, rather than an incidental direction for future work.

ACKNOWLEDGMENT

The authors thank colleagues in the Department of Physics and Engineering Physics, Obafemi Awolowo University, for helpful discussions during the preparation of this work.

➢ *Data and Code Availability*

The training scripts, raw results, and figures underlying this study are available from the corresponding author upon reasonable request. The computational environment used for the reported experiments was Python 3.12.3 with PyTorch 2.13.0 and NumPy 2.4.4 on CPU. Every reported metric traces to a checked-in script and its corresponding output file.

REFERENCES

[1]. M. A. Nielsen and I. L. Chuang, Quantum Computation and Quantum Information, 10th anniversary ed. Cambridge, U.K.: Cambridge University Press, 2010.
[2]. J. Preskill, "Quantum computing in the NISQ era and beyond," Quantum, vol. 2, p. 79, 2018, doi: 10.22331/Q-2018-08-06-79.
[3]. R. P. Feynman, "Simulating physics with computers," Int. J. Theor. Phys., vol. 21, pp. 467-488, 1982, doi: 10.1007/BF02650179.
[4]. I. M. Georgescu, S. Ashhab, and F. Nori, "Quantum simulation," Rev. Mod. Phys., vol. 86, no. 1, pp. 153-185, 2014, doi: 10.1103/RevModPhys.86.153.
[5]. G. Lindblad, "On the generators of quantum dynamical semigroups," Commun. Math. Phys., vol. 48, no. 2, pp. 119-130, 1976, doi: 10.1007/BF01608499.
[6]. H.-P. Breuer and F. Petruccione, The Theory of Open Quantum Systems. Oxford, U.K.: Oxford University Press, 2002.
[7]. H.-P. Breuer, E.-M. Laine, J. Piilo, and B. Vacchini, "Colloquium: Non-Markovian dynamics in open quantum systems," Rev. Mod. Phys., vol. 88, no. 2, p. 021002, 2016, doi: 10.1103/RevModPhys.88.021002.
[8]. S. Lloyd, "Universal quantum simulators," Science, vol. 273, no. 5278, pp. 1073-1078, 1996, doi: 10.1126/science.273.5278.1073.
[9]. D. Dong and I. R. Petersen, "Quantum control theory and applications: A survey," IET Control Theory Appl., vol. 4, no. 12, pp. 2651-2671, 2010, doi: 10.1049/iet-cta.2009.0508.

[10]. M. Raissi, P. Perdikaris, and G. E. Karniadakis, "Physics-informed neural networks: A deep learning framework for solving forward and inverse problems involving nonlinear partial differential equations," J. Comput. Phys., vol. 378, pp. 686-707, 2019, doi: 10.1016/j.jcp.2018.10.045.
[11]. G. E. Karniadakis, I. G. Kevrekidis, L. Lu, P. Perdikaris, S. Wang, and L. Yang, "Physics-informed machine learning," Nat. Rev. Phys., vol. 3, pp. 422-440, 2021, doi: 10.1038/s42254-021-00314-5.
[12]. S. Cuomo et al., "Scientific machine learning through physics-informed neural networks: Where we are and what's next," J. Sci. Comput., vol. 92, Art. no. 88, 2022, doi: 10.1007/s10915-022-01939-z.
[13]. J. Biamonte and V. Bergholm, "Tensor networks in a nutshell," arXiv:1708.00006, 2017.
[14]. G. Carleo and M. Troyer, "Solving the quantum many-body problem with artificial neural networks," Science, vol. 355, no. 6325, pp. 602-606, 2017, doi: 10.1126/science.aag2302.
[15]. G. Torlai et al., "Neural-network quantum state tomography," Nat. Phys., vol. 14, pp. 447-450, 2018, doi: 10.1038/s41567-018-0048-5.
[16]. J. Sirignano and K. Spiliopoulos, "DGM: A deep learning algorithm for solving partial differential equations," J. Comput. Phys., vol. 375, pp. 1339-1364, 2018.
[17]. J. Han, A. Jentzen, and W. E, "Solving high-dimensional partial differential equations using deep learning," Proc. Natl. Acad. Sci. USA, vol. 115, no. 34, pp. 8505-8510, 2018, doi: 10.1073/pnas.1718942115.
[18]. A. S. Olatinwo and N. A. Salako, "PINN-KK: A physics-informed neural network for causality-compliant spectroscopic phase retrieval in low-loss regimes," Int. J. Sci. Res., vol. 15, no. 8, pp. 1305-1314, Aug. 2026.
[19]. G. G. de Lima, I. Cunha, and L. K. Castelano, "Inverse physics-informed neural networks procedure for detecting noise in open quantum systems," arXiv:2507.12552, 2025.
[20]. A. Sulc, "Quantum noise tomography with physics-informed neural networks," arXiv:2509.11911, 2025.
[21]. S. Biswas and E. Paspalakis, "Learning and inverting driven open quantum systems via physics-informed neural networks," Mach. Learn.: Sci. Technol., vol. 7, no. 2, p. 025022, 2026, doi: 10.1088/2632-2153/ae4b84.
[22]. N. B. Dehaghani, A. P. Aguiar, and R. Wisniewski, "Quantum Pontryagin neural networks in Gamkrelidze form subjected to the purity of quantum channels," IEEE Control Syst. Lett., vol. 7, pp. 2227-2232, 2023, doi: 10.1109/LCSYS.2023.3286303.
[23]. N. B. Dehaghani, A. P. Aguiar, and R. Wisniewski, "A hybrid quantum-classical physics-informed neural network architecture for solving quantum optimal control problems," in Proc. IEEE Int. Conf. Quantum Comput. Eng. (QCE), 2024, doi: 10.1109/QCE60285.2024.00164.
[24]. C. P. Koch, "Controlling open quantum systems: Tools, achievements, and limitations," J. Phys.: Condens. Matter, vol. 28, no. 21, p. 213001, 2016, doi: 10.1088/0953-8984/28/21/213001.
[25]. M. Tancik et al., "Fourier features let networks learn high frequency functions in low dimensional domains," in Adv. Neural Inf. Process. Syst. (NeurIPS), vol. 33, pp. 7537-7547, 2020.
[26]. D. P. Kingma and J. Ba, "Adam: A method for stochastic optimization," in Proc. Int. Conf. Learn. Represent. (ICLR), 2015.
[27]. W. H. Press, S. A. Teukolsky, W. T. Vetterling, and B. P. Flannery, Numerical Recipes: The Art of Scientific Computing, 3rd ed. Cambridge, U.K.: Cambridge University Press, 2007.