

Solving the Quantum Many-Body Lindbladian Learning Problem With Neural Differential Equations

Roman Aseguinolaza Gallo¹

Supervised by: Timothy Heightman^{1,2}, Edward Jiang¹

¹ICFO-Institut de Ciències Fòniques, The Barcelona Institute of Science and Technology, 08860 Castelldefels (Barcelona), Spain

²Quside Technologies SL, Carrer d'Esteve Terradas, 1, 08860 Castelldefels, Barcelona, Spain

24 August 2025

Learning open quantum many-body dynamics is challenging: full Liouvillian models grow exponentially with system size, and dissipation and dephasing force us to follow mixed states from noisy, limited data. These factors make routine characterisation and control difficult, so we need methods that are data-efficient, scalable, and easy to interpret. We present an interpretable, robust framework for learning Lindbladian dynamics from minimal, hardware-friendly data. The method pairs a physics-first CPTP Lindblad model with a small Neural Differential Equation (NDE) residual and uses a two-stage curriculum (neural warm-up, then analytic-only refinement) to reliably recover coherent and dissipative parameters on challenging 1D benchmarks. There are two ways in which robustness emerges in Lindbladian learning: modest physical dissipation that smoothens loss landscapes via steady-state attraction, and the NDE residual that resolves remaining nonconvexity when paired with an optimizer reset. A transient infidelity metric shows short-time power-law error and small steady-state plateaus. Extending beyond CPTP to a stochastic dissipative qubit shows failures in noise-induced or deep PT-unbroken phases that are information-limited, not optimization-limited.

Keywords: Lindbladian Learning, Neural Ordinary Differential Equations, Quantum Many-body Systems, Quantum Machine Learning, Non-Hermitian Dynamics

Acknowledgements

I am deeply grateful to my supervisor, Tim, whose passion for research has been truly inspiring and whose guidance, freedom, and support have been invaluable throughout this work. I also wish to thank Edward and Marcin for their advice and helpful discussions. Special thanks to Quside Inc. for the compute resources.



**Master in Quantum
Science and Technology
Barcelona**

Roman Aseguinolaza Gallo¹: roman.aseguinolaza@icfo.eu

Nire guraso eta arrebari. Hemen idatzitakoa inoiz ulertuko ez duzuen arren, nik lortutakoa zuei esker da, niri bezainbeste.

Attona eta amonari. Lan honetaz harro nago, badakidalako zuek ere harro egongo zinateketelako nitaz.

Contents

1	Introduction	1
1.1	Problem Statement	2
2	Beyond-CPTP testbed. Stochastic Dissipative Qubit: master equation and phase diagram	4
3	Neural Differential Equations and Learning Framework	5
3.1	Measurements and Loss Function	6
3.2	Curriculum Learning	7
4	Experimental models and setup	7
5	Results	9
5.1	Noise impact	10
5.2	NDE-enhanced learning	11
5.3	Accuracy and Extrapolation	12
5.4	Beyond Lindblad learning: SDQ results and phase-resolved robustness . . .	14
6	Conclusion and Outlook	18
	Bibliography	19
A	GKSL Generator Invariance Under Unitary Mixing of Jump Operators	22
B	Code and Data Availability	22

1 Introduction

The reliable identification of generators governing many-body quantum dynamics is foundational for error mitigation, device certification, and model-based quantum control [1–4]. While Hamiltonian Learning (HL) has recently achieved robust performance in closed systems using Neural Differential Equations (NDEs) and experimentally friendly data protocols [5], practical platforms operate as open systems where dissipation, dephasing, and control imperfections are intrinsic [6]. This reality motivates a principled extension from unitary generators to Liouvillian learning, in which one seeks interpretable coherent terms together with physically meaningful dissipative channels from limited, noisy observations. In practice, this lets one carry out three relevant and useful tasks: tune Hamiltonian parameters even when small, unknown noise is present; measure the main noise channels when a Hamiltonian model is already trusted; or recover a complete, interpretable open-system generator using short evolutions and standard measurements. This directly supports routine calibration of noisy quantum computing hardware, and it enables control and error-mitigation workflows by recovering the dominant dissipative channels alongside the coherent terms [7–9].

The learning problem exhibits two practical formulations: the white-box and black-box scenarios. In a white-box scenario, we fix the generator’s structure in advance, and then learn only the corresponding real coefficients and nonnegative rates from data (see Fig. 1). In a black-box scenario, we do not fix the operator structure in advance: both the basis elements and their coefficients are inferred from data at the same time. This greatly enlarges the search space and introduces intrinsic non-uniqueness due to unitary gauge freedoms in the Lindblad representation (see Appendix A for a short derivation of the gauge). Gray-box methods interpolate between these extremes by fixing part of the structure (e.g., locality pattern, jump set) while learning unknown magnitudes, which is often the most realistic stance in many-body hardware where engineering and calibration coexist with latent interactions [10].




Figure 1: White-box Lindbladian learning. We assume access to time-evolved states under a CPTP dynamical map $\Phi(t) = e^{t\mathcal{L}}$ generated by a structured Lindbladian. The operator sets P_j (coherent terms) and L_α (jump operators) are fixed by prior knowledge (e.g., locality, engineering and calibration), while the unknown real parameters c_j and nonnegative rates γ_α are inferred from measurements after controllable evolution times t . Under standard locality and measurement-completeness assumptions, this parametrization yields an identifiable, finite-dimensional estimation problem.

In open-system identification, complementary strategies emerge depending on what is known a priori and which measurements are feasible. When the dissipative magnitude is weak (i.e., the norm of the generator’s dissipative part is small), short-time quenches of local systems enable joint recovery of coherent and dissipative terms with predictable shot-scaling, providing transient-time routes that avoid full tomography [11]. At the opposite temporal extreme, steady-state measurements can identify local Lindbladians—and

even aid in reconstructing commuting Hamiltonians—thereby offering a resource-efficient alternative when long-time data are available [12]. For scenarios where memory cannot be neglected, likelihood-based estimators have been proposed that reconstruct time-local generators consistent with non-Markovian dynamics, bridging toward more general dynamical-map tomography [13]. Motivated by experimental friendliness and alignment with near-term control constraints, we operate in the transient regime: prepare a small set of basis states, evolve only for short durations, and sample random Pauli observables. This setting leverages informative state trajectories while requiring no optimal control beyond choosing measurement times.

Parallel progress in machine learning indicates that neural-network based methods can distill dynamical structure from modest, noisy datasets when coupled with physical constraints and inductive biases [14]. Advances span data-driven predictors trained under random drivings, limited-data state and process tomography, expressive variational representations for mixed many-body states, and sequence models that diagnose Markovianity from time series [15–18]. These developments collectively motivate hybrid identification pipelines where minimal, interpretable physics is augmented by disciplined neural components to improve robustness and data efficiency on many-body hardware.

Building on these ideas, we extend the NDE-based HL framework from [5] to open quantum systems. We couple a physically structured Lindblad-form generator with a lightweight neural augmentation under a curriculum designed for interpretability and stable optimization. This approach retains experimental friendliness (short-time evolution, random Pauli measurements), respects completely positive and trace preserving map (CPTP) structure through the analytic component, and provides quantitative diagnostics for robustness and extrapolation in dissipative settings.

We found that adding more physical noise smoothens the loss landscape and often negates the need for NDEs. However, in realistic settings with fractionally small noise relative to the unitary operator norm, they remained necessary to be as robust as the results from unitary HL [5]. Out of sample, errors grow with a short-time power law and then flatten to small steady-state plateaus, which shows the learned fixed points are accurate. As a compact complement to our CPTP program, we also apply the same learning framework to a recently introduced stochastic non-Hermitian single-qubit model [19]. In this setting, the algorithm stays robust in regimes where the data carry enough information, and it fails in phases where fast collapse or very weak damping erase that information.

The remainder of this work is organized as follows. Section 1.1 states the problem and assumptions. Section 2 introduces the stochastic dissipative qubit and its anti-dephasing master equation. Section 3 presents the gray-box NDE framework, with measurements, loss, and the curriculum for training. Section 4 details the experimental models and data protocol. Section 5 reports results on noise-induced smoothing, NDE-enhanced learning with optimizer reset, accuracy and extrapolation, and the SDQ phase-resolved robustness. We conclude with conclusions and outlook.

1.1 Problem Statement

We consider Markovian open-system dynamics on 1D spin-1/2 chains. We assume the system plus environment evolves unitarily and the environment is not observed. Tracing out the environment, the reduced state $\rho(t)$ evolves under a time-homogeneous CPTP

semigroup generated by the Lindblad master equation [20]

$$\dot{\rho}(t) = \mathcal{L}_T[\rho(t)] \equiv -i[H_T, \rho(t)] + \sum_{\alpha} \gamma_{\alpha} \left(L_{\alpha} \rho(t) L_{\alpha}^{\dagger} - \frac{1}{2} \{L_{\alpha}^{\dagger} L_{\alpha}, \rho(t)\} \right), \quad (1)$$

with coherent Hamiltonian H_T , jump operators L_{α} , and nonnegative real-valued rates γ_{α} . The learning task is to infer a generator within a chosen model structure that explains observed data and remains physically valid (CPTP), prioritizing interpretability of both coherent and dissipative terms. By interpretability we mean: the Hamiltonian and dissipative channels are expressed in a fixed, physically motivated operator basis, and most of the fitted dynamics is attributed to these analytic terms, with any additional flexibility tightly regularized.

For spin-1/2 chains of length N , we represent the coherent part in the Pauli basis,

$$H_T = \sum_{j=0}^{4^N-1} c_j P_j, \quad P_j \in \mathcal{P}_N, c_j \in \mathbb{R}, \quad (2)$$

which is a complete decomposition for any Hamiltonian on N qubits. Dissipative channels are modeled by local jump operators (e.g., relaxation/dephasing) with rates $\gamma_{\alpha} \in \mathbb{R}_+$.

The objective is therefore to learn (c_j, γ_{α}) so that the induced trajectories match observed statistics within the training window and generalize to nearby times. The learning algorithm is based on maximum-likelihood estimation on random local Pauli measurements during transient times, avoiding full process tomography while remaining identifiable under locality assumptions [21, 22].

The measurement data is structured as follows: prepare a small set of eigenstate product states, evolve for short durations, and perform random, mutually compatible Pauli measurements with finite shots. We target sub-exponential classical post-processing with shot budgets consistent with standard quantum-limit scaling [23]. Assumptions on noise and control are minimal: we assume no state preparation and measurement errors, and no time-dependent control is applied beyond selecting evolution times.

We also assume finite-range locality (nearest-neighbour by default, with straightforward extension k -local interactions), and we operate primarily in a white-/gray-box regime (fixed locality pattern and jump set, unknown coefficients). These restrictions reduce parameter count and improve identifiability, but they can bias recovery if the true generator contains longer-range couplings or different dissipators.

When additional expressivity is required, we augment the analytic generator with a small regularized set of neural parameters acting as residual corrections, trained under a curriculum described in 3.2. This aims to mitigate mild misspecification without sacrificing interpretability.

In addition to CPTP Lindbladian learning on spin chains, we will apply the identical learning pipeline—without architectural changes—to a stochastic non-Hermitian single qubit benchmark (SDQ). This add-on exposes how our estimator behaves when trajectories are not CPTP (see Sec. 2).

2 Beyond-CPTP testbed. Stochastic Dissipative Qubit: master equation and phase diagram

We consider a system with Hermitian Hamiltonian that is subject to classical noise coupled to a non-Hermitian operator. Stochastic Hamiltonians model the environment through noise statistics, and upon noise averaging with trace renormalization one obtains a non-linear “anti-dephasing” master equation that is distinct from GKSL yet experimentally achievable [24–26].

The SDQ setting corresponds to the Lindblad evolution of a 3-level system $\{|g\rangle, |e\rangle, |f\rangle\}$, where all quantum jumps from $|e\rangle$ to $|g\rangle$ are discarded. The time-dependent non-Hermitian Hamiltonian we consider is

$$\hat{H}_t^{\text{SDQ}} = J\hat{\sigma}_x - i\Gamma(1 + \sqrt{2\gamma}\xi_t)\hat{\Pi}, \quad \hat{\Pi} = |e\rangle\langle e|, \quad (3)$$

where J is the coherent coupling between $|e\rangle$ and $|f\rangle$, $\Gamma \geq 0$ is the mean loss from $|e\rangle$, $\gamma \geq 0$ controls classical (white) noise strength, and ξ_t is Gaussian white noise [19]. The corresponding (noise-averaged, trace-renormalized) “anti-dephasing” master equation used for learning is

$$\partial_t \hat{\rho}_t = -iJ[\hat{\sigma}_x, \hat{\rho}_t] - (\Gamma - \gamma\Gamma^2)\{\hat{\Pi}, \hat{\rho}_t\} + 2\gamma\Gamma^2\hat{\Pi}\hat{\rho}_t\hat{\Pi} + (2\Gamma - 4\gamma\Gamma^2)\text{Tr}(\hat{\Pi}\hat{\rho}_t)\hat{\rho}_t. \quad (4)$$

We use this equation as-is and learn its parameters directly from transient measurements using the same NDE-based pipeline as for CPTP generators. The learnable parameters are (J, Γ, γ) . $\hat{\Pi}$ is fixed by the experimental design. Their ratios organize the dynamics into three regimes phases that are straightforward to recognize in transient signals and in the steady state, ρ_s , through $z(\rho_s) = \text{Tr}[\hat{\sigma}_z \rho_s]$. The dynamics splits into three phases: PT-unbroken (PTu), PT-broken (PTb), and a noise-induced (NI) phase. Here “PT” refers to the parity–time symmetry of the underlying non-Hermitian two-level block in the noiseless limit. Figure 2a summarizes these regimes; Figures 2b and 2c show representative transients.

PT-unbroken. Inherited from the noiseless threshold $J \geq \Gamma/2$. For small γ , the leading eigenvalues retain nonzero imaginary parts, producing visible oscillations. The dissipative gap is tiny, so relaxation is slow and the steady state lies near the maximally mixed state, $z(\rho_s) \approx 0$. This regime cleanly exposes residual damping on top of coherent oscillations, which is informative about the noise while the oscillation frequency calibrates J .

PT-broken. Dominant loss $\Gamma/J > 2$ suppresses oscillations: imaginary parts vanish deep in the phase and the dissipative gap grows. Trajectories relax quickly and monotonically toward a state close to $|f\rangle$, yielding $z(\rho_s) \rightarrow +1$. Early-time curvature and approach speed tightly constrain Γ/J .

Noise-induced. Beyond the straight critical line $\gamma^* = 1/(2\Gamma)$, classical noise stabilizes the originally leaky state $|e\rangle$: convergence becomes very fast and $z(\rho_s) \rightarrow -1$. Exactly at $\gamma = \gamma^*$ the map turns CPTP with jump $\hat{\Pi}$ at rate Γ ; away from it, rapid convergence motivates denser time sampling.

Taken together, the phase diagram and the representative transients in Fig. 2 substantiate the phase descriptions above: PTu preserves coherent oscillations with weak damping, PTb exhibits non-oscillatory relaxation set by Γ/J , and NI reveals a noise-stabilized target ($|e\rangle$) and a CPTP critical line that can serve as a calibration anchor. This structure is

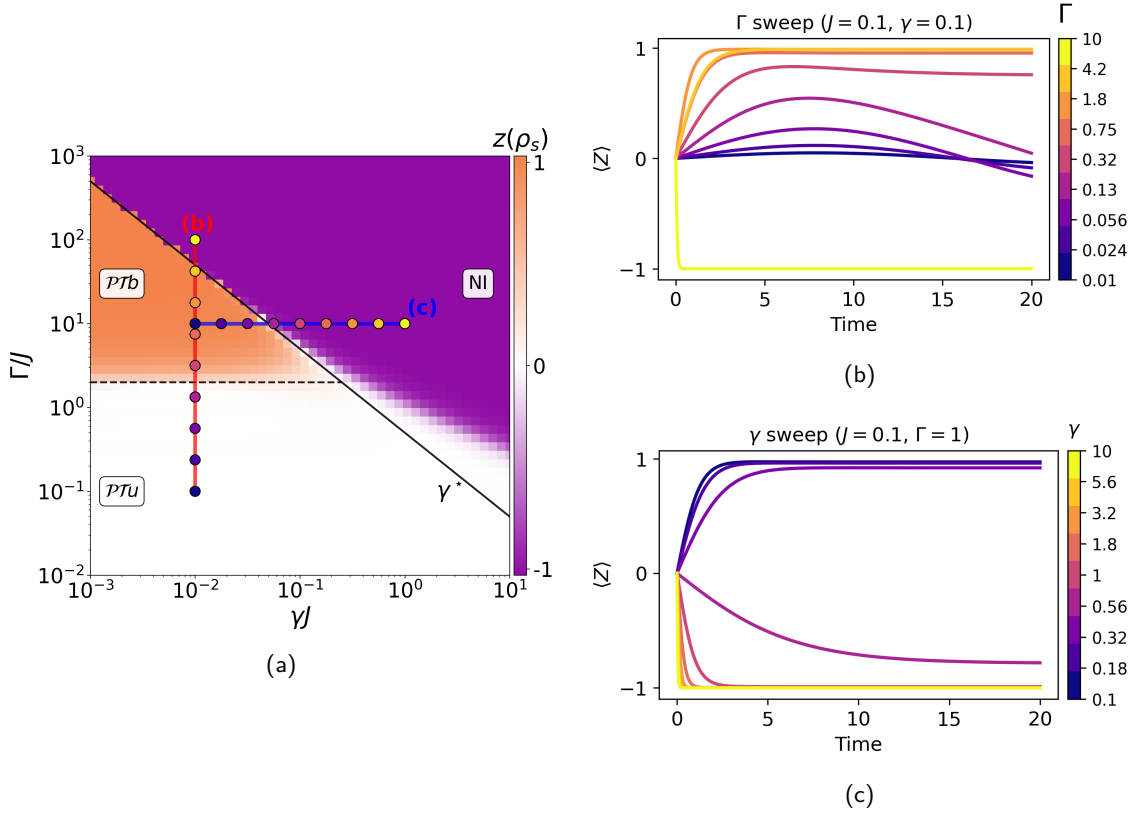


Figure 2: Stochastic dissipative qubit (SDQ). (a) Phase map in $(\Gamma/J, \gamma J)$ colored by $z(\rho_s)$ with PTu (white), PTb (orange), NI (purple). Solid diagonal: $\gamma = 1/(2\Gamma)$; dashed horizontal: $\Gamma/J = 2$. Overlaid markers indicate the Γ - and γ -sweeps used in (b,c). Initial state $\rho_0 = |+\rangle\langle+|$. (b) Transients $z(\rho_t)$ along the Γ -sweep at fixed J, γ : oscillatory and slow in PTu, fast monotone in PTb, accelerated near NI. (c) Transients along the γ -sweep at fixed J, Γ : crossing γ^* stabilizes $|e\rangle$ (negative z) with rapid NI convergence. At γ^* the dynamics is CPTP with jump $\hat{\Pi}$ at rate Γ .

interesting for control because noise becomes a tunable resource. Such phase-aware tuning suggests practical avenues for state preparation and stabilization with fluctuating gain/loss [27].

3 Neural Differential Equations and Learning Framework

A Neural Ordinary Differential Equation (NODE) specifies the vector field of an ODE with a trainable function. Given a state vector $x(t) \in \mathbb{R}^d$ and a neural architecture $f_\theta : \mathbb{R}^d \rightarrow \mathbb{R}^d$ with parameters θ , the NODE dynamics

$$\frac{dx}{dt}(t) = f_\theta(x(t)), \quad x(0) = x_0, \quad (5)$$

are integrated to time T , and the parameters θ are tuned by backpropagating a task loss through the ODE solver. Universal approximation results justify the expressivity of f_θ , and practical training uses stochastic optimizers such as Adam [14, 28, 29].

In our open-system setting the dynamical variable is a density operator $\rho(t)$ evolving under a Liouvillian superoperator. We adopt a gray-box decomposition that preserves

interpretability while permitting data-driven corrections:

$$\mathcal{L}_\Theta = \mathcal{L}_{\text{phys}}(\theta) + \mathcal{L}_{\text{NN}}(\phi), \quad (6)$$

where the physics ansatz is the standard Lindblad generator

$$\mathcal{L}_{\text{phys}}(\theta)[\rho] = -i[H(\theta_H), \rho] + \sum_{\alpha} \gamma_{\alpha}(\theta_{\gamma_{\alpha}}) \left(L_{\alpha} \rho L_{\alpha}^{\dagger} - \frac{1}{2} \{L_{\alpha}^{\dagger} L_{\alpha}, \rho\} \right). \quad (7)$$

Here θ_H are the Hamiltonian Pauli coefficients and $\theta_{\gamma_{\alpha}}$ are real-valued free parameters mapped to nonnegative rates by a softplus transformation,

$$\gamma_{\alpha}(\theta_{\gamma}) = \text{softplus}(\theta_{\gamma_{\alpha}}) = \log(1 + e^{\theta_{\gamma_{\alpha}}}), \quad (8)$$

which guarantees $\gamma_{\alpha} \geq 0$ and improves sensitivity of gradients for small rates. The neural correction $\mathcal{L}_{\text{NN}}(\phi)$ is implemented as a small, trainable superoperator acting on ρ whose structure is chosen to preserve trace and Hermiticity. Its magnitude is controlled by strong L_2 regularization and by the training curriculum so that the analytic term (7) remains the dominant term in the inferred dynamics for maximal interpretability. We emphasize that we do not impose additional physics-informed penalties beyond the softplus transformation for dissipative rates.

The master equation then reads

$$\dot{\rho}(t) = (\mathcal{L}_{\text{phys}}(\theta) + \mathcal{L}_{\text{NN}}(\phi))[\rho(t)], \quad \rho(0) = \rho_0, \quad (9)$$

and is integrated with an automatic-differentiation-compatible ODE solver. In practice we use a high-order explicit Runge–Kutta integrator with adaptive step control (Diffax/JAX) [14, 30].

3.1 Measurements and Loss Function

To train the NDE, we require a loss function that compares an estimated time-evolved state to observed data in an experimentally friendly manner. Access to the “ground truth” evolution is assumed only through samples obtained from compatible Pauli measurements on the true output state; full state or process tomography is not used.

Let $\rho_{\Theta}(t)$ denote the model-predicted state at time t obtained by integrating Eq. (9) with parameters $\Theta = (\theta, \phi)$ from a specified initial state $\rho(0)$. For each chosen Pauli measurement basis, we collect binary bitstrings $b \in \{0, 1\}^n$, forming a dataset D of measurement outcomes. Following [31], with J timestamps, M repeats per measurement, and K random Pauli bases, one has $|D| = MKJ$ samples for a single initial state. Varying the number of initial states multiplies the dataset size correspondingly.

Since our pipeline is numerical, we have full access to $\rho_{\Theta}(t)$ and can evaluate Born-rule probabilities for observing bitstrings b in the chosen Pauli basis,

$$p_{\Theta}(b; t) = \text{Tr}(\Pi_b \rho_{\Theta}(t)), \quad (10)$$

where Π_b is the projector associated with b in that Pauli basis. The training objective is the average negative log-likelihood over the observed outcomes,

$$L(\rho_{\Theta}; D) = - \sum_{b \in D'} \log p_{\Theta}(b; t), \quad (11)$$

where $D' \subseteq D$ denotes those outcomes acquired from evolving the corresponding initial state and measuring at the same timestamp used to evaluate $p_\Theta(b; t)$. This loss is differentiable with respect to Θ through the ODE solver, enabling end-to-end gradient-based updates of the analytic parameters θ and the neural parameters ϕ .

As the neural component can be substantially more expressive than the constrained Lindbladian ansatz, there is a risk that the network overfits and overwhelms the interpretable part. Since our aim in the white-/gray-box setting is an interpretable estimate—where as much of the learned dynamics as possible is attributed to the analytic Lindbladian—we must bias training toward reliance on the physics model. We do so with strong regularization of the neural parameters [32], and, crucially, with a curriculum that stages the optimization to favor the analytic component [33, 34].

3.2 Curriculum Learning

We employ a simple, effective curriculum that gradually reduces reliance on the neural augmentation while refining the interpretable parameters:

Warm-up Jointly train the analytic parameters θ and the neural parameters ϕ . This phase leverages the smoothing effect of the neural term to escape poor basins and accelerate convergence.

Analytic refinement Switch off the neural term (both forward and backward), reset the optimizer state, and train only θ . The reset is essential to restore sensitivity to small dissipative rates parameterized via the softplus map, ensuring accurate recovery of the interpretable Hamiltonian and Lindbladian coefficients.

Optional residual fine-tuning With θ frozen, one may briefly train ϕ to capture small systematic discrepancies not represented in the analytic ansatz. In our synthetic studies—where the ground truth lies in the analytic class—we omit this step.

This staged procedure mirrors the rationale behind curriculum learning: begin with a more permissive model to navigate rugged landscapes, then constrain capacity to distill an interpretable solution. In practice, the warm-up removes sensitivity to initialization and barren plateau effects, while the refinement phase consolidates the physically meaningful parameters.

4 Experimental models and setup

This section specifies the dynamical models and the data protocol used throughout. We study two many-body Hamiltonians under local Lindblad noise—the anisotropic Heisenberg chain and the constrained PXP model—and, separately, the single-qubit non-Hermitian testbed (SDQ) described in 2 while reusing the same learning pipeline.

We consider 1D spin- $\frac{1}{2}$ chains with $N \in \{3, \dots, 7\}$ qubits, where the reduced state $\rho(t)$ evolves under a time-homogeneous Lindblad master equation $\dot{\rho}(t) = \mathcal{L}[\rho(t)]$ with coherent Hamiltonian H and local jump operators. Following previous benchmarks for challenging unitary dynamics [5], we select two models that probe complementary regimes of identifiability and dynamical complexity: (i) the anisotropic Heisenberg chain,

$$H_H = \sum_{i=1}^{N-1} (J_i^x X_i X_{i+1} + J_i^y Y_i Y_{i+1} + J_i^z Z_i Z_{i+1}) + \sum_{i=1}^N h_i X_i, \quad (12)$$

and (ii) the PXP model [35], a paradigmatic weakly non-ergodic Hamiltonian,

$$H_{\text{PXP}} = \sum_{i=2}^{N-1} J_i P_{i-1} X_i P_{i+1}, \quad \text{with} \quad P_i = |0\rangle\langle 0|_i. \quad (13)$$

Both were previously identified as hard instances for closed-system HL due to rugged loss landscapes and non-ergodic dynamics, respectively, serving as stringent tests for robustness and generalization of our open-system learning algorithm [5]. To model realistic dissipation relevant to near-term hardware, we couple each site to a thermal-like environment using local ladder jump operators defined as

$$L_i^- = \sigma_i^- = \frac{1}{2}(\sigma_i^x - i\sigma_i^y), \quad L_i^+ = \sigma_i^+ = \frac{1}{2}(\sigma_i^x + i\sigma_i^y), \quad (14)$$

with nonnegative rates $\gamma_i^-, \gamma_i^+ \geq 0$, implementing amplitude damping and thermal excitation. This choice captures dominant T_1 -type relaxation and finite-temperature re-excitation processes in many platforms (e.g., superconducting circuits and Rydberg arrays) [36, 37], while preserving complete positivity and trace preservation by construction within the GKSL form. Thermal-like channels also induce unique steady states that regularize long-time dynamics [38, 39], improving identifiability of both coherent and dissipative parameters from short-time transients and stabilizing extrapolation beyond the training window. These properties align with steady-state identifiability results and transient-time learning strategies reported in recent open-system learning literature [11, 40].

Ground-truth datasets are generated by numerically integrating the Lindblad master equation with the specified H and jump set L_i^\pm , using an adaptive high-order explicit integrator applied directly to the density matrix representation. Unless stated otherwise, Hamiltonian parameters are drawn i.i.d. from $[-1, 1]$, and dissipative rates are set by sampling γ_i^-, γ_i^+ in $[0.2, 1.0]$ and rescaling by a global noise level η to probe robustness across weak-to-moderate dissipation. For each $N \in \{3, 4, 5, 6, 7\}$ we prepare $L = 5$ product eigenstates of single-qubit Pauli operators, drawn uniformly from the sets of X -, Y - and Z -eigenstates (i.e., “up”/“down” along x, y, z), and evolve to $J = 10$ timestamps $t \in \{0.1, 0.2, \dots, 1.0\}$ (in units where $\hbar = 1$). Measurements follow the experimentally friendly protocol from [5]: at each (state, t) pair, we choose $K = 100$ random compatible Pauli bases and collect $M = 100$ shots per basis. Training uses the average negative log-likelihood under the Born rule computed from predicted $\rho(t)$ and the observed bitstrings. This transient, random-Pauli scheme is consistent with identifiability from local measurements and avoids oracle access or steady-state-only requirements.

For each model and system size, we generate 50 independent ground-truth parameter sets θ . Let $\tilde{\theta}$ denote the learned estimator after training via the chosen learning algorithm (vanilla or NDE-enhanced). We calculate the relative absolute error for the coherent parameters,

$$\epsilon_H(\tilde{\theta}) = \frac{|\theta_H - \tilde{\theta}_H|_1}{|\theta_H|_1}, \quad (15)$$

and, when Lindbladian parameters are learned, the analogous error for the dissipative parameters,

$$\epsilon_L(\tilde{\theta}) = \frac{|\theta_L - \tilde{\theta}_L|_1}{|\theta_L|_1}. \quad (16)$$

A trial is declared successful for parameters $S \in H, L$ if $\epsilon_S(\tilde{\theta}) < 0.1$. This threshold mirrors prior practice and cleanly separates converged from failed runs across all experiments.

Reported success rates are the proportion of successful trials over the 50 draws for fixed (H, N) and fixed experimental condition (e.g., noise level, learning algorithm).

To quantify generalization, we evaluate the density-matrix fidelity

$$F(\rho_T(t), \rho(t; \theta)) = \left(\text{Tr} \sqrt{\sqrt{\rho_T(t)} \rho(t; \theta) \sqrt{\rho_T(t)}} \right)^2 \quad (17)$$

and define the infidelity loss

$$\mathcal{I}(\rho_T(t), \rho(t; \theta)) = 1 - F(\rho_T(t), \rho(t; \theta)). \quad (18)$$

5 Results

We evaluate the framework on 1D spin-1/2 chains with local Lindblad noise, and then on the SDQ testbed. The aim is to determine when the analytic Lindblad model suffices, when a small neural residual is needed, and how well the learned dynamics extrapolate beyond the training window.

We begin with robustness on spin chains. First, we study how physical dissipation alone improves convergence in a hard, non-ergodic setting (PXP) by smoothing the loss landscape. Next, we turn to a higher-parameter model (anisotropic Heisenberg) and show that a short neural warm-up followed by an optimizer reset removes residual nonconvexity and recovers both Hamiltonian and Lindbladian parameters reliably. We then assess out-of-sample behaviour using the infidelity benchmark. Finally, we apply the same pipeline to the SDQ testbed to probe behaviour beyond CPTP. This separates optimisation effects from limits imposed by the data.

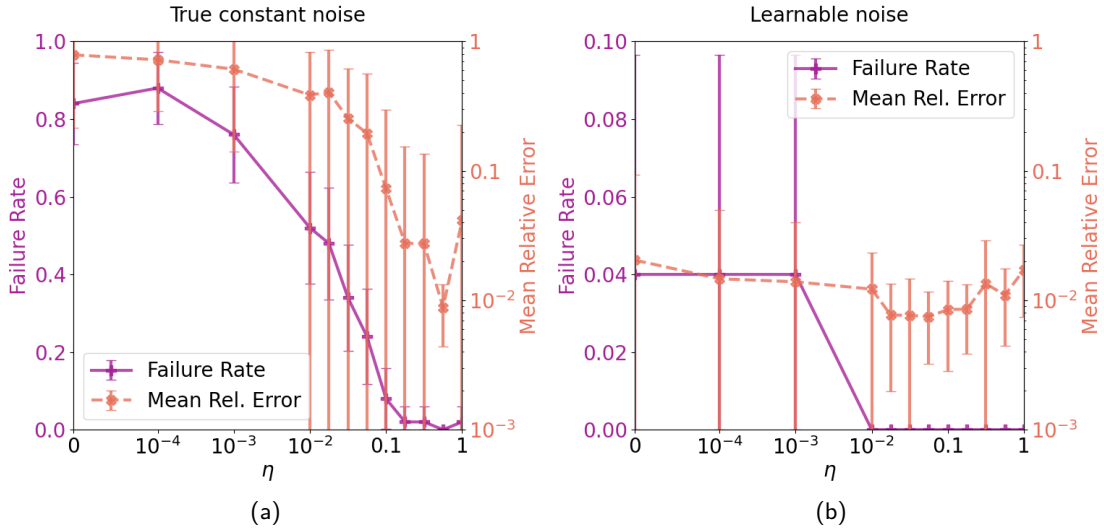


Figure 3: Impact of injected thermal noise level η on white-box learning for PXP at $N = 3$. Each panel shows (purple) failure rate and (orange, log-scale) mean relative error of the Hamiltonian coefficients versus η . (a) dissipative rates are fixed to their ground-truth values and not learned. Increasing noise monotonically reduces failures, illustrating that strong dissipation regularizes the Hamiltonian estimation. Failures drop below 10% once rates are ~ 0.1 of coherent scales. (b) dissipative rates are learned jointly with the Hamiltonian. During training, the rates transiently rise before converging to their true values, which induces the same robustness mechanism and yields accurate Hamiltonian recovery even when the true noise is negligible.

5.1 Noise impact

We probe how thermal noise impacts convergence in a setting designed to stress identifiability in the unitary limit: the PXP Hamiltonian, whose constrained, non-ergodic dynamics yield rugged landscapes and high failure rates for vanilla white-box learning at zero dissipation (see Fig. 3a). To isolate the role of dissipation, we start with the smallest nontrivial instance, $N = 3$, where we can scan noise settings comprehensively and still capture the qualitative behavior that persists at larger sizes.

We sweep a global noise level $\eta \in [0, 1]$ that multiplicatively rescales all dissipative rates, keeping the sampling protocol otherwise unchanged. We first freeze the dissipative rates to their ground-truth values and attempt to learn only the Hamiltonian parameters using the vanilla learning algorithm. Close to the noiseless limit, convergence becomes fragile, which is typical for constrained, non-ergodic dynamics. While weak dissipation improves convergence and stabilizes training to some extent, robust performance requires a moderate to substantial noise level. Optimization remains fragile near the noiseless limit, and only stronger dissipation ensures consistently reliable convergence.

Figure 3 illustrates the measured success-rate transition as a function of η , while Fig. 4 shows representative two-dimensional slices of the loss landscape with and without dissipation. The landscapes reveal that dissipation “rounds” sharp valleys and reduces the prevalence of local minima, producing gradient fields that guide optimization toward the global minima even from typical random initializations. This smoothing effect is qualitatively similar to the neural augmentation observed in closed-system Hamiltonian learning (see [5] Fig. 3), both mechanisms reduce ruggedness and shrink the set of local minima. However, the smoothing here arises from physical mixing and the emergence of an attractive steady state that is common across initial conditions. Intuitively, the relaxation pathways bridge dynamically separated regions that, in the purely constrained, non-ergodic limit, would trap local search. This makes for easier convergence from generic initializations.




Figure 4: Loss landscapes for the two PXP Hamiltonian parameters at system size $N = 3$ under increasing thermal noise η . Each panel is centered at the global minimum and axes are the two physical parameters of the PXP model. As η grows, the landscape transitions from (a) highly rugged with many spurious minima and saddle features to (c) a markedly smoother, bowl-shaped basin around the optimum. Thus, dissipation acts as a landscape regularizer, enlarging the region of attraction of the global minimum and reducing the chance of getting trapped away from it. This smoothing mechanism is consistent with the robustness trends in Fig. 3: higher noise lowers failure rates and improves Hamiltonian recovery.

Table 1: Robustness on the PXP Hamiltonian with thermal-like damping. Reported are success rates (in %) over 50 runs with random parameter initializations, learned with engineered noise and no NDE. Top block: Hamiltonian learning, with no true noise (true $\eta = 0$). Bottom block: joint recovery of Hamiltonian and dissipator (H, L) with learnable rates, initialized at $0.1\times$ the coherent scale.

Noise	N	Success % (H)	Success % (L)
0	3	96	–
0	4	98	–
0	5	98	–
$0.1\times$	3	100	100
$0.1\times$	4	100	100
$0.1\times$	5	100	100

Allowing the algorithm to estimate the Lindbladian rates (rather than freezing them to ground truth) preserves the same qualitative trends. Even when the true system has vanishingly small dissipation, introducing a weak learnable dissipative sector significantly reduces failure rates for the coherent parameters (see Fig. 3b). In low-noise regimes, failures fall to $\sim 4\%$, and for $\eta \gtrsim 0.01$ we observe near-perfect success in simultaneous H -only identification with fixed Lindblad, while full (H, L) identification reaches 100% in the same range. Table 1 summarizes the trend for $N \in \{3, 4, 5\}$.

Taken together, these results motivate an “engineered-noise” perspective: in challenging non-ergodic instances, controlled weak dissipation can act as a regularizer that widens basins of attraction and enhances learning. In practice, modest thermalization is ubiquitous on near-term platforms, and these observations suggest that appropriately accounting for it can turn a liability into an advantage for learning. While learning with neural augmentations can also smooth landscapes, dissipation offers a physics-grounded route that is architecture-agnostic and complements neural methods. We emphasize that engineered-noise learning is not a substitute for neural augmentation on higher-parameter Hamiltonians; rather, it provides a baseline improvement in robustness and, when combined with NDEs, yields consistently high success rates across models.

5.2 NDE-enhanced learning

We next evaluate the impact of neural differential equation augmentation in the anisotropic Heisenberg model, whose parameter space is substantially higher dimensional than that of the PXP and exhibits more complex loss geometries. Here, the vanilla algorithm already benefits from moderate dissipation but continues to show degradation at intermediate sizes (e.g., $N = 4, 6$), reflecting residual landscape complexity and sensitivity to initialization. Table 2 reports baseline vanilla success rates with thermal-like damping and no NDE.

Introducing the NDE augmentation with the standard two-stage curriculum (warm-up with NN on, then NN off, train analytic terms) improves coherent-parameter robustness but initially causes a pronounced trade-off (see Table 2): Lindbladian parameters degrade very substantially.

This trade-off arises because, during the warm-up, the NN model is able to represent the noise very well at transient times. Small-rate dissipative gradients are comparatively weak, so when the NN is disabled without resetting the optimizer, the stale adaptive moments

Table 2: Effect of NDE augmentation with curriculum and optimizer reset on anisotropic Heisenberg with thermal Lindblad noise. Success rates (in %) over 50 trials per size. Side-by-side comparison shows that the optimizer reset removes the earlier trade-off, yielding near-perfect joint recovery.

N	Success % (H)			Success % (L)		
	Vanilla	NDE	NDE + reset	Vanilla	NDE	NDE + reset
3	98	98	98	98	82	98
4	92	98	98	92	32	98
5	98	-	100	98	-	100
6	92	-	100	90	-	100
7	98	-	100	98	-	98

downweight precisely those directions now responsible for fitting the residual dynamics. The softplus-reparametrized rates then stall near the noise floor.

An optimizer-state reset at the curriculum boundary resolves this problem. By clearing accumulated moments when switching off the NN, we restore sensitivity to the dissipative-rate gradients, enabling accurate joint refinement of H and L . With this single change, the NDE-enhanced learning algorithm achieves near-perfect success across sizes, eliminating the earlier trade-off. Table 2 shows the resulting success rates.

The robustness picture that emerges is cohesive across models and mechanisms. In non-ergodic cases like PXP, a small amount of dissipation already makes training reliable. For higher-parameter models like the anisotropic Heisenberg chain, dissipation helps but is not enough on its own; a short neural warm-up is needed to get past the remaining optimisation hurdles. That warm-up, however, can leave the optimiser biased against small dissipative rates. Resetting the optimiser at the curriculum switch fixes this and lets both the Hamiltonian and the Lindbladian be recovered accurately. In practice, combining weak physical dissipation (present on realistic devices) with NDE augmentation and an optimizer reset yields high success rates in both coherent and dissipative sectors, while keeping most of the learned dynamics in the interpretable analytic model.

5.3 Accuracy and Extrapolation

We evaluate accuracy and out-of-sample behavior with the infidelity benchmark introduced for Hamiltonian Learning in [5]. In the unitary case, the error accumulated by the learned flow is well described by a power law in time: for ground-truth and predicted pure states $|\psi(t)\rangle$ and $|\psi'(t)\rangle$, the infidelity scales as

$$\mathcal{I}(t) = 1 - F(|\psi(t)\rangle, |\psi'(t)\rangle) = At^b, \quad (19)$$

with coefficients (A, b) characteristic of the algorithm and problem class. It provides a model-agnostic way to track local error growth beyond the training window time. In open quantum systems, we extend the same idea by fitting a power law only on the transient portion of the dynamics, now using density-matrix fidelity introduced in 4.

Consistent with the unitary benchmark, very short-time behavior exhibits a power-law regime that reflects the global integration error accumulated by the learned flow. Unlike the closed-system setting, however, the presence of thermal jump operators drives trajectories to steady states, so the power law does not persist indefinitely: after an initial rise,




Figure 5: Infidelity for anisotropic Heisenberg dynamics with thermal Lindblad noise. For each system size $N \in \{3, 4, 5, 6, 7\}$, the curves show the average over 30 random Hamiltonian instances and 10 random initial states (empirical mean across all runs). The shaded region marks the training-time window, times to the right are out-of-sample. Dashed lines indicate short-time power-law fits used only to read off the pairs (A_N, b_N) . The horizontal red dashed line marks the 1% error guide.

$\mathcal{I}(t)$ reaches a maximum in the transient window and subsequently decays toward a constant plateau. This asymptotic value quantifies the mismatch between the learned and true steady states and provides a natural open-system analogue of long-horizon generalisation.

Figure 5 shows $\mathcal{I}(t)$ for anisotropic Heisenberg Hamiltonians with thermal Lindbladians at $N \in \{3, 4, 5, 6, 7\}$. The shaded region marks the training-time window; dashed lines indicate short-time power-law fits used solely to extract exponents and prefactors. Two qualitative features are robust across sizes: (i) transient dynamics are harder to reproduce than steady-state behavior—infidelity increases within the first decade beyond the training window before relaxing; and (ii) the final plateau is small, indicating that the learned generator captures the dominant dissipative fixed point. For the system sizes considered, peak infidelities remain well below 10^{-2} , highlighting strong extrapolation to unseen times.

To characterise the short-time scaling, we fit $\mathcal{I}(t) \approx A_N t^{b_N}$ over an interval strictly contained in the pre-asymptotic regime, avoiding both the training window and the onset of steady-state relaxation. Beyond (A_N, b_N) , two additional open-system benchmarks are informative: the peak infidelity in the transient regime and the long-time plateau (steady-state mismatch). Table 3 below reports all four quantities with standard deviations where applicable.

The fitted short-time power laws, with prefactors A_N and exponents b_N , align closely with those reported for Hamiltonian-only learning (see [5] Table 2), indicating that extending the protocol to open dynamics preserves performance under this metric rather than degrading it. In particular, b_N varies mildly with N and falls within the same narrow ranges observed in the closed-system study [5]. This points to b_N being a property of the learning pipeline (data protocol + optimizer + model class), not

Table 3: Short-time power-law fits $\mathcal{I}(t) \approx A_N t^{b_N}$ and open-system benchmarks (peak and steady-state infidelities) for $N \in \{3, 4, 5, 6, 7\}$.

(a) Fit coefficients (mean \pm s.d.).		
N	A_N	b_N
3	$2.272 \times 10^{-4} \pm 9.469 \times 10^{-7}$	1.288 ± 0.016
4	$3.083 \times 10^{-4} \pm 1.359 \times 10^{-6}$	1.275 ± 0.016
5	$4.563 \times 10^{-4} \pm 2.066 \times 10^{-6}$	1.197 ± 0.017
6	$4.821 \times 10^{-4} \pm 4.453 \times 10^{-6}$	1.299 ± 0.035
7	$5.280 \times 10^{-4} \pm 7.300 \times 10^{-6}$	1.395 ± 0.052

(b) Infidelity summary.		
N	Max infidelity	Final (steady-state) infidelity
3	1.099×10^{-3}	4.596×10^{-5}
4	1.093×10^{-3}	1.017×10^{-4}
5	1.390×10^{-3}	1.079×10^{-4}
6	1.380×10^{-3}	2.151×10^{-4}
7	1.772×10^{-3}	7.455×10^{-4}

of any specific ground-truth example, and makes b_N a handy indicator of out-of-sample performance for Lindbladian learning. The weak dependence of b_N on N also suggests the protocol can scale, provided the underlying simulation cost is controlled. Finally, this benchmark is easy to reuse with any Lindbladian-learning method: given a learned generator, evolve some held-out initial states under both the true and learned dynamics beyond the training window and fit the log-log slope of infidelity to get b .

The good long-time behavior is naturally explained by the same effect seen throughout the paper—moderate dissipation smooths the loss landscape and stabilizes training—so the algorithm tracks the dominant dissipative fixed point while still reproducing the relevant transients. The subsequent decay to a small, size-dependent plateau reflects convergence toward a learned steady state close to the true fixed point; the plateau height offers a concise metric of steady-state accuracy complementary to the short-time power-law benchmark. That said, we emphasize a limitation: our present experiments cover $N \leq 7$, and it remains an open question how the peak-and-plateau behavior evolve at larger N , where state-space growth makes systematic studies significantly more challenging.

5.4 Beyond Lindblad learning: SDQ results and phase-resolved robustness

We close the results with a non-Hermitian benchmark that lies outside the CPTP setting analyzed so far. We keep the same gray-box learning pipeline and curriculum used for spin chains, and adapt only the data protocol to accommodate the extreme relaxation scales that arise across SDQ phases.

We draw parameter triples (J, Γ, γ) log-uniformly from $[0.03, 3.0]^3$ and then augment the set to obtain 50 instances per dynamical phase—PT-unbroken (PTu), PT-broken (PTb), and noise-induced (NI)—for a total of 150 ground-truth models. We will evaluate robustness on recovering the coherent parameter, J , and the parameters related to the non-Hermitian operator, Γ and γ , separately.




Figure 6: Vanilla algorithm on the SDQ: phase-resolved outcomes overlaid on the steady-state map $z(\rho_s)$. Markers encode recovery success by sector: both H and L (green), H (J) only (blue), L (Γ, γ) only (orange), both failed (red). The horizontal dashed line $\Gamma/J = 2$ marks the PT threshold; the solid diagonal $\gamma^* = 1/(2\Gamma)$ is the CPTP line. Robust recovery concentrates in PT-broken; failures cluster deep in PT-unbroken (weak damping) and in NI (rapid collapse).

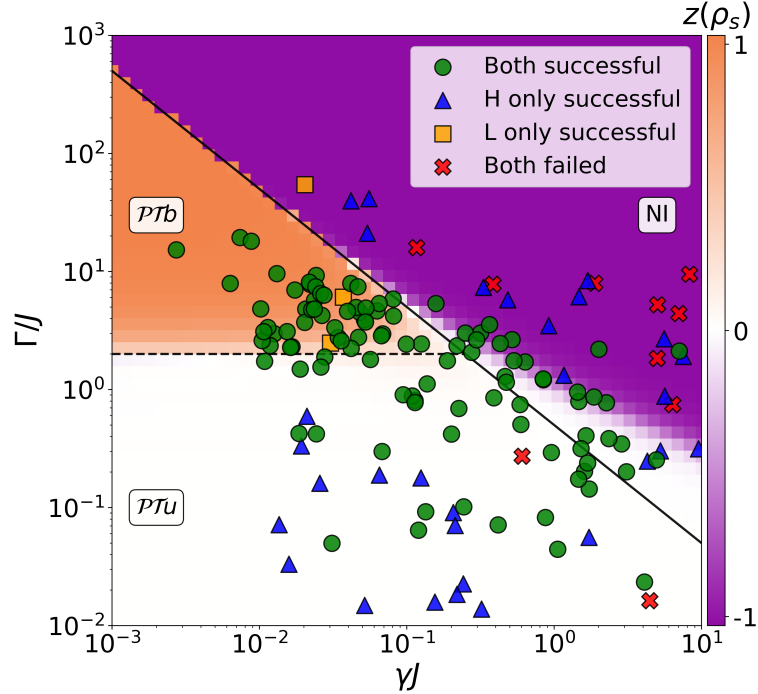


Figure 7: NDE-augmented algorithm (warm-up + optimizer reset) on the SDQ, same operating points and legend as Fig. 6. Near the CPTP line $\gamma^* = 1/(2\Gamma)$, both $H(J)$ and $L(\Gamma, \gamma)$ are reliably recovered. Outside PT-broken and away from the CPTP corridor, added expressivity yields little benefit—the limiting factor is data informativeness under fast NI collapse or too-weak PT-unbroken damping.

Table 4: Phase-resolved robustness on the SDQ. Columns report success rates (%) for coherent ($H:J$) and dissipative ($L:\Gamma, \gamma$) parameters under three protocols: NDE with coarse internal step ($\Delta t = 0.1$), NDE with dense internal step ($\Delta t = 0.01$), and vanilla ($\Delta t = 0.01$). Dense stepping improves L recovery where informative transients exist (near CPTP and in PT-broken), while both methods degrade in NI and deep PT-unbroken due to limited information rather than optimization.

Phase	$\Delta t = 0.1$ neural		$\Delta t = 0.01$ neural		$\Delta t = 0.01$ Vanilla	
	H %	L %	H %	L %	H %	L %
PTu	88.5	46.1	96	66	98	68
PTb	85.7	100	96	100	96	98
NI	90	40	82	54	88	60
TOTAL	88	60	91.3	73.3	94	75.3

To counteract low sensitivity of noise parameters in extreme regimes, We collect denser training data by sampling measurements every $\Delta t = 0.01$ and we extend training to 100 epochs. We compare the vanilla analytic algorithm against the NDE-augmented one with a warm-up and optimizer reset at the curriculum boundary.

The phase-resolved maps in Figs. 6–7 and the summary in Table 4 reveal a clear dependence on dynamical regime. As expected, in PT-broken the vanilla algorithm is already very robust: relaxation is monotone with a sizable dissipative gap, so both H (J) and L (Γ, γ) parameters are identified reliably. Fast entry into NI and modest excursions into PT-unbroken degrade learning of the noise parameters in particular. The mechanism differs from the many-body (spin-chain) failures addressed by NDEs: here the bottleneck is information, not optimization. In NI, trajectories collapse extremely rapidly to the same steady state across a wide swathe of (Γ, γ) , suppressing transient structure and erasing Fisher information needed to disentangle rates. In deep PT-unbroken, oscillations pin J accurately but damping is too slow to resolve small noise within feasible windows, again limiting information. Because the data lack sensitivity in these regimes, extra expressivity from NDEs yields little to no improvement away from the CPTP corridor. We note that in the PT-unbroken and NI phases, but near the CPTP line, $\gamma^* = 1/(2\Gamma)$, convergence slows and parameters are learned with ease.

As an explicit illustration of information-limited regimes in the SDQ, we compute the time-resolved classical Fisher information (CFI) for the noise parameter γ at representative operating points across the phase diagram. The Fisher information quantifies the information about the parameter γ that can be extracted from observation of a probabilistic observable. For classical projective measurements, the relevant quantity is the CFI,

$$F_C(t; \gamma) = \sum_j \frac{1}{p_j(t)} \left(\frac{\partial p_j(t)}{\partial \gamma} \right)^2, \quad (20)$$

where $p_j(t)$ are the Born-rule probabilities of outcomes m_j at time t for a chosen initial state [41]. In our single-qubit readout, $p_j(t) = \text{Tr}(\Pi_j, \rho_t)$ with Π_j the projectors of X , Y , or Z , and ρ_t the solution of Eq. (4). Figure 8 summarizes $F_C(t; \gamma)$ along a γ -sweep at fixed (J, Γ) for six initial states $|0\rangle, |1\rangle, |+\rangle, |-\rangle, |+i\rangle, |-i\rangle$ and the three Pauli bases.

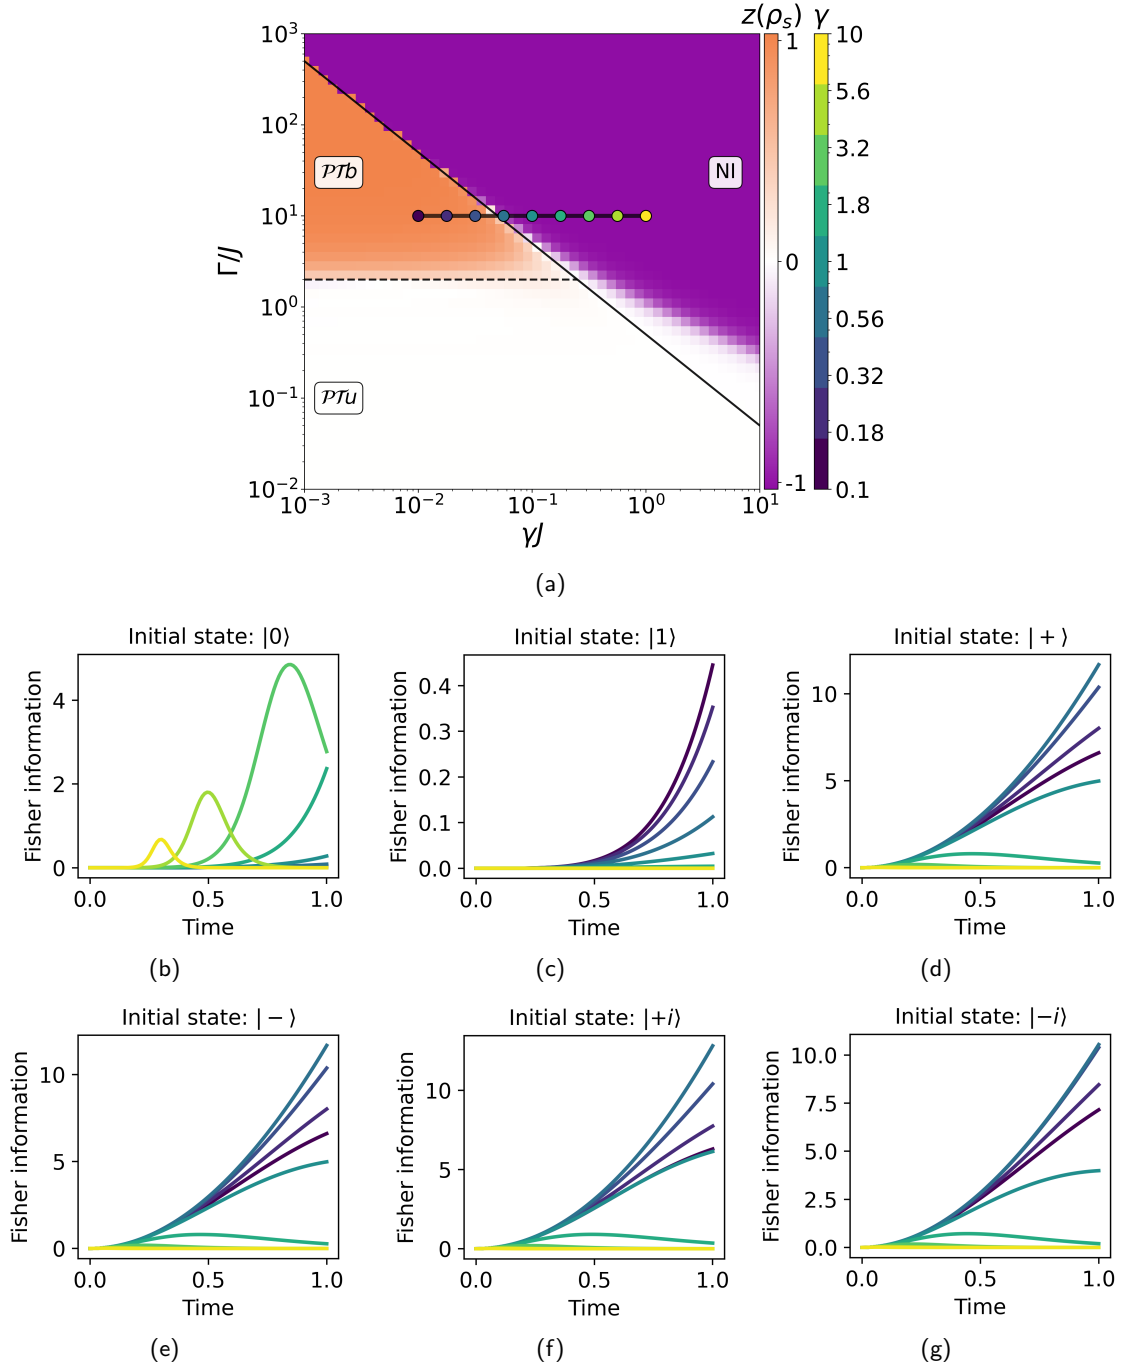


Figure 8: Gamma-sweep Fisher information in the stochastic dissipative qubit. (a) Phase map in $(\Gamma/J, \gamma J)$ as in Fig. 2, with the γ -sweep operating points overlaid. (b–g) Time-resolved Fisher information $F_C(t; \gamma)$ available from single-qubit Pauli measurements X , Y , and Z (the sum of the 3 is given) for six initial states $|0\rangle$, $|1\rangle$, $|+\rangle$, $|-\rangle$, $|+i\rangle$, $| -i\rangle$. Each panel shows $F_C(t; \gamma)$ for several γ along the sweep at fixed $(J = 0.1, \Gamma = 1)$, revealing phase-dependent informativeness.

Two phase-dependent trends are apparent. First, for most initial states the points in the PT-broken phase provide substantially larger and more broadly distributed information about γ than points deep in the NI phase: $F_C(t; \gamma)$ grows monotonically, enabling reliable estimation. Second, in NI the information collapses rapidly. $F_C(t; \gamma)$ concentrates into a narrow early-time peak whose height diminishes as one moves deeper into NI, reflecting

the swift attraction to a common steady state that washes out parameter dependence. The initial state $|0\rangle$ is a useful edge case: while it retains some sensitivity to γ , the deeper the operating point lies in NI, the sharper and lower the $F_C(t; \gamma)$ peak becomes, leaving only a short temporal window with modest information content. These CFI profiles are consistent with the phase-resolved learning outcomes in Figs. 6–7 and Table 4: parameter recovery is robust where F_C is large over a finite time interval (PT-broken and near the CPTP line), and it degrades when F_C is either too small (deep PT-unbroken) or too localized and rapidly extinguished (deep NI).

Training with sparser external sampling ($\Delta t = 0.1$) underperforms across phases (bottom block of Table 4). This degradation directly corroborates the sensitivity explanation: with fewer, more widely spaced snapshots, the short-lived or weak dissipative signatures that identify γ and its interplay with Γ are not sufficiently expressed in the data.

Overall, the SDQ study complements the spin-chain analysis: NDEs are most valuable when failures originate from rugged landscapes. When failures are driven by fast attraction (NI) or too-weak damping (deep PTu), the limiting factor is data informativeness. In those regimes, experimental design (earlier/denser sampling or calibrated dissipation) is the appropriate lever, while the analytic estimator already achieves phase-aware robustness near the CPTP corridor and throughout PTb.

6 Conclusion and Outlook

This work extends Neural Differential Equations from closed to open quantum dynamics by pairing a physically structured Lindblad generator with a lightweight neural residual trained under a simple two-stage curriculum. In an experimentally friendly setting—short evolutions and random Pauli measurements—the method reliably recovers both coherent and dissipative parameters under locality assumptions, stabilizing optimization in many-body settings with rugged, nonconvex loss landscapes. Altogether, the pipeline remains interpretable and CPTP-consistent through its analytic core, while the neural augmentation compensates mild misspecification in a controlled way.

From an application standpoint, the framework supports white- and gray-box workflows, and even black-box workflows if the NDE were to dominate over the ansatz. Experimentalists seeking to identify a system’s Hamiltonian in the presence of weak, unknown noise should prioritize a white-/gray-box fit of the analytic Lindblad model. Those wishing to perform dominant noise-channel characterization under a trusted Hamiltonian should freeze H and train only the dissipator within the analytic class, inserting a brief neural warm-up and an optimizer reset if convergence stalls. Under full Liouvillian learning of $(H, \{L_k\})$, use the same short neural warm-up and optimizer reset as above, but do not freeze any parameters. These workflows map directly to quantum simulation, error mitigation, and device certification by enabling data-driven open-system models with controlled interpretability, calibration and compensation of coherent and incoherent drift, and likelihood-based verification of controlled operations under a structured generator.

We also tested the pipeline beyond CPTP on the stochastic dissipative qubit (SDQ) governed by a nonlinear anti-dephasing master equation. Identifiability is phase dependent: it is strong in PT-broken, degrades in NI due to rapid collapse to a common steady state, and weakens deep in PT-unbroken where damping is too small to resolve. Robust performance reappears near the CPTP line $\gamma^* = 1/(2\Gamma)$, where transients carry information.

Several limitations point to concrete next steps. Scalability is the central bottleneck: despite locality and short-time protocols, density-matrix propagation grows exponentially with system size. Incorporating tensor-network or reduced-order methods for Lindbladians, along with structure-exploiting integrators, is a natural way to expand the accessible system sizes while preserving differentiability and CPTP structure. Another route is to work in multi-qubit phase-space: represent states and dynamics as smooth functions over a domain that scales linearly with qubit number, instead of exponentially large Hilbert-space objects. This function-space view enables variational modelling directly on phase-space. For technical details and proofs, see the recent multi-qubit phase-space foundations [42].

On the data side, random Pauli marginals are hardware-friendly but can bias likelihood estimates under finite shots and SPAM drifts. Integrating shadow-based estimators and related variance-reduction techniques into the learning objective could improve sample efficiency without altering the basic workflow (e.g., classical shadows and shadow tomography provide principled ways to estimate many observables from few measurements [43]). For non-Hermitian stochastic dissipative qubit, scalability is not the bottleneck; instead, data expressivity dominates, which also suggests allocating budget to denser, phase-aware sampling, or other means of information gathering altogether, rather than to larger models or solvers.

Bibliography

- [1] Armands Strikis, Dayue Qin, Yanzhu Chen, Simon C. Benjamin, and Ying Li. “Learning-based quantum error mitigation”. *PRX Quantum* **2** (2021).
- [2] Elizabeth R. Bennewitz, Florian Hopfmueller, Bohdan Kulchytskyy, Juan Carrasquilla, and Pooya Ronagh. “Neural error mitigation of near-term quantum simulations”. *Nature Machine Intelligence* **4**, 618–624 (2022).
- [3] Jens Eisert, Dominik Hangleiter, Nathan Walk, Ingo Roth, Damian Markham, Rhea Parekh, Ulysse Chabaud, and Elham Kashefi. “Quantum certification and benchmarking”. *Nature Reviews Physics* **2**, 382–390 (2020).
- [4] Irtaza Khalid, Carrie A. Weidner, Edmond A. Jonckheere, Sophie G. Schirmer, and Frank C. Langbein. “Sample-efficient model-based reinforcement learning for quantum control”. *Phys. Rev. Res.* **5**, 043002 (2023).
- [5] Timothy Heightman, Edward Jiang, and Antonio Acín. “Solving the quantum many-body hamiltonian learning problem with neural differential equations” (2024). [arXiv:2408.08639](https://arxiv.org/abs/2408.08639).
- [6] John Preskill. “Quantum computing in the nisq era and beyond”. *Quantum* **2**, 79 (2018).
- [7] Piotr Czarnik, Andrew Arrasmith, Patrick J. Coles, and Lukasz Cincio. “Error mitigation with clifford quantum-circuit data”. *Quantum* **5**, 592 (2021).
- [8] Armands Strikis, Dayue Qin, Yanzhu Chen, Simon C. Benjamin, and Ying Li. “Learning-based quantum error mitigation”. *PRX Quantum* **2**, 040330 (2021).
- [9] Gabriel O. Samach, Ami Greene, Johannes Borregaard, Matthias Christandl, Joseph Barreto, David K. Kim, Christopher M. McNally, Alexander Melville, Bethany M. Niedzielski, Youngkyu Sung, Danna Rosenberg, Mollie E. Schwartz, Jonilyn L. Yoder, Terry P. Orlando, Joel I-Jan Wang, Simon Gustavsson, Morten Kjaergaard, and William D. Oliver. “Lindblad tomography of a superconducting quantum processor”. *Physical Review Applied* **18** (2022).
- [10] Akram Youssry, Yang Yang, Robert J. Chapman, Ben Haylock, Francesco Lenzini,

- Mirko Lobino, and Alberto Peruzzo. “Experimental graybox quantum system identification and control”. *npj Quantum Information* **10**, 9 (2024).
- [11] Tobias Olsacher, Tristan Kraft, Christian Kokail, Barbara Kraus, and Peter Zoller. “Hamiltonian and liouvillian learning in weakly-dissipative quantum many-body systems”. *Quantum Science and Technology* **10**, 015065 (2025).
- [12] Eyal Bairey, Chu Guo, Dario Poletti, Netanel H Lindner, and Itai Arad. “Learning the dynamics of open quantum systems from their steady states”. *New Journal of Physics* **22**, 032001 (2020).
- [13] S. Varona, M. Müller, and A. Bermudez. “Lindblad-like quantum tomography for non-markovian quantum dynamical maps”. *npj Quantum Information* **11**, 96 (2025).
- [14] Patrick Kidger. “On neural differential equations” (2022). [arXiv:2202.02435](#).
- [15] Naeimeh Mohseni, Thomas Fösel, Lingzhen Guo, Carlos Navarrete-Benlloch, and Florian Marquardt. “Deep Learning of Quantum Many-Body Dynamics via Random Driving”. *Quantum* **6**, 714 (2022).
- [16] Akshay Gaikwad, Omkar Bihani, Arvind, and Kavita Dorai. “Neural-network-assisted quantum state and process tomography using limited data sets”. *Phys. Rev. A* **109**, 012402 (2024).
- [17] Hannah Lange, Anka Van de Walle, Atiye Abedinnia, and Annabelle Bohrdt. “From architectures to applications: a review of neural quantum states”. *Quantum Science and Technology* **9**, 040501 (2024).
- [18] Angela Rosy Morgillo, Massimiliano F. Sacchi, and Chiara Macchiavello. “Detecting markovianity of quantum processes via recurrent neural networks”. *Quantum Machine Intelligence* **7**, 47 (2025).
- [19] Pablo Martinez-Azcona, Aritra Kundu, Avadh Saxena, Adolfo del Campo, and Aurélia Chenu. “Quantum dynamics with stochastic non-hermitian hamiltonians”. *Physical Review Letters* **135** (2025).
- [20] G. Lindblad. “On the generators of quantum dynamical semigroups”. *Communications in Mathematical Physics* **48**, 119–130 (1976).
- [21] Daniel Burgarth and Kazuya Yuasa. “Identifiability of open quantum systems”. *Phys. Rev. A* **89**, 030302 (2014).
- [22] Robbie King, David Gosset, Robin Kothari, and Ryan Babbush. “Triply efficient shadow tomography”. *PRX Quantum* **6**, 010336 (2025).
- [23] Vittorio Giovannetti, Seth Lloyd, and Lorenzo Maccone. “Quantum-enhanced measurements: Beating the standard quantum limit”. *Science* **306**, 1330–1336 (2004).
- [24] M. Naghiloo, M. Abbasi, Yogesh N. Joglekar, and K. W. Murch. “Quantum state tomography across the exceptional point in a single dissipative qubit”. *Nature Physics* **15**, 1232–1236 (2019).
- [25] Yu-Liang Fang, Jun-Long Zhao, Yu Zhang, Dong-Xu Chen, Qi-Cheng Wu, Yan-Hui Zhou, Chui-Ping Yang, and Franco Nori. “Experimental demonstration of coherence flow in pt- and anti-pt-symmetric systems”. *Communications Physics* **4** (2021).
- [26] Alexander Quinn, Jeremy Metzner, Jacob E. Muldoon, Isam Daniel Moore, Sean Brudney, Sourin Das, David T. C. Allcock, and Yogesh N. Joglekar. “Observing super-quantum correlations across the exceptional point in a single, two-level trapped ion” (2023). [arXiv:2304.12413](#).
- [27] Frank Verstraete, Michael M. Wolf, and J. Ignacio Cirac. “Quantum computation and quantum-state engineering driven by dissipation”. *Nature Physics* **5**, 633–636 (2009).
- [28] Ricky T. Q. Chen, Yulia Rubanova, Jesse Bettencourt, and David K Duvenaud. “Neural ordinary differential equations”. In S. Bengio, H. Wallach, H. Larochelle, K. Grauman, N. Cesa-Bianchi, and R. Garnett, editors, *Advances*

- in Neural Information Processing Systems. Volume 31. Curran Associates, Inc. (2018). url: https://proceedings.neurips.cc/paper_files/paper/2018/file/69386f6bb1dfed68692a24c8686939b9-Paper.pdf.
- [29] Diederik P. Kingma and Jimmy Ba. “Adam: A method for stochastic optimization” (2017). [arXiv:1412.6980](#).
 - [30] John Charles Butcher and Gerhard Wanner. “Runge-kutta methods: some historical notes”. *Applied Numerical Mathematics* **22**, 113–151 (1996).
 - [31] Frederik Wilde, Augustine Kshetrimayum, Ingo Roth, Dominik Hangleiter, Ryan Sweke, and Jens Eisert. “Scalably learning quantum many-body hamiltonians from dynamical data” (2022). [arXiv:2209.14328](#).
 - [32] Federico Girosi, Michael Jones, and Tomaso Poggio. “Regularization theory and neural networks architectures”. *Neural Computation* **7**, 219–269 (1995).
 - [33] Yoshua Bengio, Jérôme Louradour, Ronan Collobert, and Jason Weston. “Curriculum learning”. In *Proceedings of the 26th Annual International Conference on Machine Learning*. Page 41–48. ICML ’09New York, NY, USA (2009). Association for Computing Machinery.
 - [34] Petru Soviany, Radu Tudor Ionescu, Paolo Rota, and Nicu Sebe. “Curriculum learning: A survey”. *Int. J. Comput. Vision* **130**, 1526–1565 (2022).
 - [35] Maksym Serbyn, Dmitry A. Abanin, and Zlatko Papić. “Quantum many-body scars and weak breaking of ergodicity”. *Nature Physics* **17**, 675–685 (2021).
 - [36] P. Krantz, M. Kjaergaard, F. Yan, T. P. Orlando, S. Gustavsson, and W. D. Oliver. “A quantum engineer’s guide to superconducting qubits”. *Applied Physics Reviews* **6** (2019).
 - [37] Bang Liu, Li-Hua Zhang, Yu Ma, Qi-Feng Wang, Tian-Yu Han, Jun Zhang, Zheng-Yuan Zhang, Shi-Yao Shao, Qing Li, Han-Chao Chen, Guang-Can Guo, Dong-Sheng Ding, and Bao-Sen Shi. “Bifurcation of time crystals in driven and dissipative rydberg atomic gas”. *Nature Communications* **16**, 1419 (2025).
 - [38] E. B. Davies. “Markovian master equations”. *Communications in Mathematical Physics* **39**, 91–110 (1974).
 - [39] E. B. Davies. “Markovian master equations. ii”. *Mathematische Annalen* **219**, 147–158 (1976).
 - [40] Chi-Sheng Chen and En-Jui Kuo. “Unraveling quantum environments: Transformer-assisted learning in lindblad dynamics” (2025). [arXiv:2505.06928](#).
 - [41] Michael A. Nielsen and Isaac L. Chuang. “Quantum computation and quantum information: 10th anniversary edition”. Cambridge University Press. Cambridge (2010).
 - [42] Timothy Heightman, Edward Jiang, Ruth Mora-Soto, Maciej Lewenstein, and Marcin Płodzień. “Quantum machine learning in multi-qubit phase-space part i: Foundations” (2025). [arXiv:2507.12117](#).
 - [43] Jonathan Kunjummen, Minh C. Tran, Daniel Carney, and Jacob M. Taylor. “Shadow process tomography of quantum channels”. *Phys. Rev. A* **107**, 042403 (2023).

A GKSL Generator Invariance Under Unitary Mixing of Jump Operators

Proof. We recall the Gorini–Kossakowski–Sudarshan–Lindblad (GKSL) master equation for a density operator ρ ,

$$\dot{\rho} = -i[H, \rho] + \mathcal{D}(\rho), \quad \mathcal{D}(\rho) = \sum_k \gamma_k \left(L_k \rho L_k^\dagger - \frac{1}{2} \{L_k^\dagger L_k, \rho\} \right), \quad (21)$$

where H is Hermitian, L_k are jump (collapse) operators, and the rates satisfy $\gamma_k \in \mathbb{R}_+$.

Consider the transformation of the collapse operators as

$$\sqrt{\gamma_k} L_k \rightarrow \sum_j U_{kj} \sqrt{\gamma_j} L_j,$$

where U is a unitary matrix, i.e., $\sum_k U_{kj} U_{kl}^* = \delta_{jl}$.

Under this transformation, the dissipator transforms like

$$\begin{aligned} \mathcal{D}(\rho) &\rightarrow \sum_k \left(\sum_j U_{kj} \sqrt{\gamma_j} L_j \rho \sum_l U_{kl}^* \sqrt{\gamma_l} L_l^\dagger - \frac{1}{2} \left\{ \sum_j U_{kj}^* \sqrt{\gamma_j} L_j^\dagger \sum_l U_{kl} \sqrt{\gamma_l} L_l, \rho \right\} \right) \\ &= \sum_{j,l} \left(\sum_k U_{kj} U_{kl}^* \sqrt{\gamma_j \gamma_l} L_j \rho L_l^\dagger - \frac{1}{2} \left\{ \sum_k U_{kj}^* U_{kl} \sqrt{\gamma_j \gamma_l} L_j^\dagger L_l, \rho \right\} \right) \\ &= \sum_{j,l} \left(\delta_{jl} \sqrt{\gamma_j \gamma_l} L_j \rho L_l^\dagger - \frac{1}{2} \left\{ \delta_{jl} \sqrt{\gamma_j \gamma_l} L_j^\dagger L_l, \rho \right\} \right) \\ &= \sum_j \gamma_j \left(L_j \rho L_j^\dagger - \frac{1}{2} \{L_j^\dagger L_j, \rho\} \right) \\ &= \mathcal{D}(\rho). \end{aligned}$$

Hence, the GKSL generator is invariant under any unitary rotation among the jump operators $\sqrt{\gamma_k} L_k$, and the corresponding dynamics are unchanged. \blacksquare

B Code and Data Availability

The full source code implementing the learning pipeline will be released in a public repository concurrent with the posting of this research to the arXiv. Prior to public release the repository may be shared on reasonable request to the author for academic assessment purposes.