Learning Volterra Kernels for Non-Markovian Open Quantum Systems

Jimmie Adriazola[1] and Katarzyna Roszak[2]

[1] *School of Mathematical and Statistical Sciences, Arizona State University, Tempe, USA*
[2] *FZU - Institute of Physics of the Czech Academy of Sciences, 182 00 Prague, Czech Republic*

(Dated: January 14, 2026)

We develop a data-driven framework for identifying non-Markovian dynamical equations of motion for open quantum systems. Starting from the Nakajima–Zwanzig formalism, we vectorize the reduced density matrix into a four-dimensional state vector and cast the dynamics as a Volterra integro-differential equation with an operator-valued memory kernel. The learning task is then formulated as a constrained optimization problem over the admissible operator space, where correlation functions are approximated by rational functions using Padé approximants. We establish well-posedness of the learning problem, ensuring existence of minimizers. To assess performance, we construct synthetic data sets from representative test problems, including (i) a solvable model with correlation functions expressed in terms of special functions, (ii) the full quantum dot model with frequency-resolved bath integrals, and (iii) a non-block-diagonal kernel arising from noncommuting system–bath couplings with cross-correlations. Numerical experiments demonstrate that Padé capture nontrivial temporal structures such as oscillatory memory, algebraic tails, and phase-sensitive coherence transfer, and that the learned models generalize across ensembles of physically admissible initial states. Our results illustrate that data-driven rational approximation provides an effective route to identifying non-Markovian kernels of practical relevance in quantum technologies.

I. INTRODUCTION

The rapid development of quantum technologies has renewed attention to the accurate modeling of open quantum systems, where a system of interest interacts with an uncontrolled environment. Quantum bits (qubits), such as quantum dots, or molecular excitons all suffer from decoherence and dissipation arising from these couplings, limiting performance in quantum information processing, nanoscale devices, and spectroscopy. The standard Markovian approximation, which leads to Lindblad or Gorini–Kossakowski–Sudarshan–Lindblad (GKSL) master equations, is often inadequate in regimes where system–bath correlations persist on time scales comparable to the system dynamics [1–3]. Capturing such *non-Markovian* features remains a central challenge for both theory and simulation.

A principled starting point is the Nakajima–Zwanzig projection operator formalism [4, 5], which yields a Volterra integro–differential equation for the reduced density matrix with an operator–valued memory kernel. Analytically, the kernel can be expressed in terms of bath correlation functions, but these are rarely tractable beyond toy models or Gaussian environments. This has led to substantial efforts across multiple communities to *learn* memory kernels from data. In molecular dynamics and statistical mechanics, generalized Langevin equations (GLEs) with memory kernels have been identified from trajectory data using projection techniques, kernel regression, and rational function approximations [6–11]. Machine learning methods, ranging from Gaussian processes to deep neural networks, have also been applied to infer nonlocal kernels in coarse–grained models of polymers and biomolecules [12, 13]. These works demonstrate the feasibility of extracting Volterra–type dynamics directly from time–series data.

In the quantum setting, recent attention has turned to learning Markovian and non-Markovian master equations directly. Approaches include fitting time–nonlocal kernels from process tomography [14], reconstructing transfer tensors [15], and more recently applying machine learning to Lindblad generators with state–dependent corrections [16] and within the so-called GENERIC formalism [17]. These methods highlight both the promise and the difficulty: one must balance expressiveness of the learned kernel with physical constraints such as complete positivity, trace preservation, and Hermiticity.

Our contribution in this work is indebted to the methods developed by these communities. Here, we synthesize various tools to develop a bare bones, data–driven framework that uses rational approximants (Padé) to represent bath correlation functions with a forward evaluation of the dynamics handled by a non-local Crank-Nicolson method. Crucially, via a Tikhonov regularization, we ensure stable Volterra kernels while retaining flexibility to capture the dynamic features of open quantum systems.

Beyond methodology, the physical *relevance* of accurate kernel identification is underscored by several model problems. First, exactly solvable toy models such as a two–level system coupled to a bosonic bath yield correlation functions expressible in terms of special functions [18, 19], providing valuable benchmarks for algorithms. Second, semiconductor quantum dots coupled to phonon baths are experimentally realized testbeds where frequency–resolved bath integrals determine decoherence times [20, 21], making them central to solid–state quantum technologies. Third, cases with noncommuting system–bath couplings (e.g. simultaneous σ_x and σ_z interactions) produce non–block–diagonal memory kernels with cross–correlations [22, 23], challenging both theory and numerics and demanding new learning strategies. By designing synthetic data around these representative sce-

narios, we demonstrate that Padé parameterizations can faithfully recover nontrivial temporal structures including oscillatory memory, algebraic tails, and coherence transfer.

This paper establishes a direct bridge between applied mathematics and quantum engineering by formulating the identification of non-Markovian kernels as a constrained optimization problem over operator spaces. Our analysis establishes existence of minimizers, connects to classical Volterra kernel learning, and validates numerical performance on physically relevant testbeds. In doing so, we provide both a mathematical framework and practical tools for quantum technologies where non-Markovianity cannot be ignored. Due to the simplicity of our approach, we aim to reveal as clearly as possible the numerical challenges ahead with learning matrix-valued Volterra kernels in open quantum settings.

The paper is organized as follows. Sec. II describes the methodology used. Secs III, IV, and V apply the methodology to different physical problems with consecutively growing complexity. Sec. VI concludes the paper.

II. OVERALL METHODOLOGY

The theoretical study of open quantum systems is complicated by the enormous dimensionality of the full system–environment Hilbert space. In practice, one is rarely interested in the environment degrees of freedom explicitly, but rather in the properties of the subsystem of interest. These can be described by the reduced density matrix

$$\rho(t) \;=\; \mathrm{Tr}_E\,\sigma(t),$$

obtained by tracing out the environment from the total system–environment density operator $\sigma(t)$. This construction dramatically reduces the effective state space, yet retains all the physically relevant information: the reduced density matrix yields the correct statistics for any observable acting on the subsystem alone, rendering the full $\sigma(t)$ unnecessary for the prediction of measurable outcomes [1–3]. Consequently, the analysis of open quantum dynamics is typically carried out at the level of $\rho(t)$ rather than the joint system–environment state.

A fundamental difficulty is that the evolution of $\rho(t)$ is no longer unitary, since environmental degrees of freedom act as an effective reservoir of dissipation and noise. In special situations, the reduced dynamics can be solved exactly. One canonical example is the spin–boson model, which describes a two-level system linearly coupled to a bath of harmonic oscillators. This model serves as a testbed for decoherence, dissipation, and quantum phase transitions, and admits closed-form solutions for certain parameter regimes [18, 19]. However, such analytically tractable cases are rare.

In more general settings, perturbative methods are applied. The most common is the Born approximation, which assumes weak system–environment coupling (and is equivalent to second order perturbative expansion with respect to the coupling) and leads to master equations of Gorini–Kossakowski–Sudarshan–Lindblad (GKSL) type under an additional Markov approximation [1, 24, 25]. These models have been commonly used in many scenarios, including quantum optics and solid state qubits, but they fail in regimes where system–bath correlations persist on timescales comparable to system evolution. Capturing such features requires either going beyond the Markov approximation or taking into account higher order processes, and motivates the development of alternative frameworks, such as projection operator techniques [4, 5, 26–29], transfer tensor methods [15], and data-driven approaches.

Often, one resorts to the Born approximation in order to obtain tractable models of qubits coupled to large and structured environments [1], but, for example, in the paradigmatic spin–boson model [18, 19], the second-order approximation is not sufficient to capture the full evolution of the spin qubit in many parameter ranges [korbicz]. The reduced qubit density matrix $\rho_S(t)$ for this problem satisfies a non-local time equation of the form

$$i\frac{d\rho_S(t)}{dt} = [H_S, \rho_S(t)] - i\int_0^t \Big([\sigma_z,\, \sigma_z(\tau - t)\rho_S(\tau)]\, C(t-\tau)\Big)\, d\tau. \tag{1}$$

Here H_S is the Hamiltonian of the system, σ_z is the appropriate Pauli matrix, and $C(t-\tau)$ denotes the bath correlation function associated with the oscillator environment. This structure illustrates the essential feature: the subsystem dynamics acquire memory terms that integrate over the entire history of the qubit–bath interaction.

More generally, the Nakajima–Zwanzig equation for a qubit can be written in Volterra integro–differential form [4, 5, 26]:

$$\frac{d\rho}{dt} = \mathcal{L}\rho + \int_0^t \mathcal{K}(t-\tau)\,\rho(\tau)\,d\tau, \tag{2}$$

where $\mathcal{L}$ is the local Liouvillian generating local unitary dynamics, and $\mathcal{K}$ is the memory kernel encoding the influence of the environment. The analytic structure of $\mathcal{K}$ is determined by bath correlation functions, but explicit expressions are typically unavailable beyond Gaussian or perturbative settings [2, 3]. Note, that for the form of eq. (2) too be sufficient to describe a physical scenario, two assumptions have to hold [1]. One is that odd moments of the interaction Hamiltonian are zero when taken with respect to the initial state of the environment, and the other that the initial system-environment state has product form. Both assumptions are commonly fulfilled.

To recast eq. (2) in a form more amenable to numerical discretization and data–driven inference, we vectorize the 2×2 qubit density matrix ρ into a four–dimensional state

vector

$$\mathbf{x}(t) \;=\; \big[\rho_{00}(t),\ \rho_{11}(t),\ \rho_{01}(t),\ \rho_{10}(t)\big]^{\intercal}.$$

In these coordinates the dynamics reduce to a linear Volterra system

$$\frac{d\mathbf{x}}{dt} = A\mathbf{x}(t) + \int_0^t B(t-\tau)\,\mathbf{x}(\tau)\,d\tau, \tag{3}$$

with $A \in \mathbb{C}^{4\times 4}$ representing the instantaneous generator and $B(\cdot)$ a matrix–valued kernel encoding time–nonlocal correlations. Equation (3) is the basic form we adopt for our learning problem.

The central objective of this work is to identify A and B directly from data. For A, the hypothesis class is straightforward: it is a fixed 4×4 complex matrix, requiring the determination of 16 parameters. For B, the hypothesis class is substantially richer: each entry $B_{ij}(t)$ is a time–dependent correlation function whose structure must be captured in a way that is both numerically tractable and physically consistent.

In the examples following this section, it will become clear that a symbolic library of elementary transcendental functions will not suffice to capture the behavior of the correlation kernel. At the same time, we do not seek to overparametrize the problem or use difficult to interpret architectures such as deep neural networks. In this work, we propose to use the $[p/q]$ Pade approximant

$$\tilde{B}_{i,j}^{[q/r]}(t;\xi) := \frac{\sum_{k=0}^{q}\xi_k t^k}{\sum_{k=q+1}^{q+r+1}\xi_k t^{k-q-1}}. \tag{4}$$

to model the behavior of each correlation function, where the explicit parametrization is given by the vector $\xi \in \mathbb{C}^{q+r+2}$. Assuming each entry in the correlation kernel has the same order $[q/r]$, the parametrization tensor of the correlation kernel is given by $\Xi \in \mathbb{C}^{4\times4\times(q+r+2)}$. Together with the 32 real parameters needed to learn A, the total number of real parameters is $N_\xi = 32(q+r+2)+32$.Before moving forward, we remark that the use of a Pade approximant at this stage is a computational choice whose viability will be demonstrated throughout this paper.

With these hypotheses in hand, we define the admissible operator space that we propose to search over. First, recall that the Sobolev space, $H^s([0,T])$, is a Hilbert space defined as

$$H^s([0,T]) := \left\{ f \in L^2([0,T]) \;:\; \frac{d^k f}{dt^k} \in L^2([0,T]) \right\},$$

with corresponding norm

$$\|f\|^2_{H^s([0,T])} := \sum_{k=0}^{s} \left\| \frac{d^k f}{dt^k} \right\|^2_{L^2([0,T])},$$

for all integers $0 \le k \le s$. Then, the proposed operator hypothesis space for our data-driven learning problem is given in compact form as

$$\mathcal{O} = \left\{ \begin{matrix} \tilde{A} \in \mathbb{C}^{4\times4}, \\ \tilde{B} \in H^1\big([0,T];\mathbb{C}^{4\times4}\big) \;\mid\; \tilde{B}_{ij}(t;\xi) \text{ satisfies (4)} \end{matrix} \right\} \tag{5}$$

where the Bochner space

$$H^s\left(\Omega;\mathbb{C}^{4\times4}\right) := \left\{ B : \Omega \to \mathbb{C}^{4\times4} \mid B_{ij} \in H^s(\Omega)\ \forall\ i,j \right\}$$

is a Sobolev space of functions, defined on the Borel measurable set Ω, taking values in $\mathbb{C}^{4\times4}$. No further structure on the matrices $\tilde{A}$ and $\tilde{B}$ is assumed at this level, but will be enforced and discussed in the context of the example problems throughout this paper.

We now define the objective functional that models the data-driven discovery of Nakajima-Zwanzig equations. Our first contribution to the objective functional is the loss function

$$\mathcal{J}_{\text{loss}}[A,B] = \sum_j \int_0^T \left| x_j^{\text{data}}(t) - x_j^{\text{learned}}(t) \right|^2 dt \tag{6}$$

where $\mathbf{x}^{\text{data}}$ is a time-series of data, either synthetic or experimental, while $\mathbf{x}^{\text{learned}}$ is the learned dynamics satisfying eq. (3) corresponding to the search over the operator space $\mathcal{O}$.

The second contribution to the objective functional is a weighted Tikhonov regularization, balancing an $L^2([0,T])$ penalization and an $L^2([0,T])$ penalization on the time derivatives. A regularization is necessary given the ill-posed nature of learning kernels that satisfy a Volterra equation of the first kind. The regularization we used is expressed as

$$\mathcal{J}_{\text{reg}}[B] = \sum_{i,j} (1-\beta) \left\| \tilde{B}_{i,j} \right\|^2_{L^2([0,T])} + \beta \left\| \frac{d\tilde{B}_{i,j}}{dt} \right\|^2_{L^2([0,T])}. \tag{7}$$

We note that, theoretically, for the problem to remain well-posed in the space $\mathcal{O}$, it is sufficient to penalize with the $H^1\left([0,T];\mathbb{C}^{4\times4}\right)$ norm; see Appendix B. However, to obtain desirable numerical results, we found it advantageous to weigh these two norms on $L^2([0,T])$ with an empirically chosen weighting parameter $\beta \in (0,1)$.

Thus, with a hypothesis space and an objective functional, we can now pose the learning of non-Markovian open quantum systems modeled by the Nakajima-Zwanzig equation (2) as the constrained optimization problem

$$\min_{\{A,B\}\in\mathcal{O}} \mathcal{J}[A,B] = \min_{\{A,B\}\in\mathcal{O}} (1-\alpha)\mathcal{J}_{\text{loss}}[A,B] + \alpha\mathcal{J}_{\text{reg}}[B] \tag{8}$$

subject to eq. (3), where $\alpha \in (0,1)$ is a parameter that balances between operator discovery and smoothness of the discovered operators. Given our explicit parametrization, the search is performed over the parameter vector $\xi \in \mathbb{R}^{N_\xi}$, is done so using a variant of the BFGS method

through MATLAB's `fminunc`. We discuss the computational choices made in further designing and implementing this abstract regression problem in the following examples.

III. TEST PROBLEM 1: SPIN-BOSON MODEL: PURE DEPHASING

Pure dephasing (or pure decoherence) [30–32] occurs when the interaction between a system and an environment does not involve the exchange of energy, but is related to the transfer of information about the system state to the environment [30]. This means that the environment cannot affect the occupations of the system in the so called pointer basis (the basis that diagonalizes both the system Hamiltonian and the interaction Hamiltonian), while it causes a decay of the off-diagonal elements of the density matrix in this basis. If the system is a qubit, then the environment-induced evolution is limited to a single element of the density matrix (the coherence). Thus pure dephasing of the qubit is described by the simplest version of Volterra integro-differential equation, eq. (2).

The spin-boson model [33–36] is one of the canonical examples of a model that leads to qubit pure dephasing. The Hamiltonian is given by

$$H = \frac{\varepsilon}{2}\sigma_z + \sum_k \hbar\omega_k b_k^\dagger b_k + \sigma_z \sum_k \left(f_k b_k^\dagger + f_k^* b_k\right). \quad (9)$$

Here, the first term describes the free evolution of the qubit (spin), the second term is the free Hamiltonian of the bosonic environment, and the third term describes the interaction. The operators $b_k^\dagger$ and b_k are bosonic creation and annihilation operators corresponding to wavevector k, respectively, and the corresponding energies are given by $\hbar\omega_k$. In the interaction, f_k denote the coupling constants, while ε is the energy difference between the two spin states.

This model is exactly solvable [1] and for the environment initially in a Gibbs state corresponding to temperature T, the interaction-picture evolution of the qubit coherence is given by an equation of the form

$$\frac{d\rho_{01}(t)}{dt} = -\int_0^t \rho_{01}(s)C(t-s)ds. \quad (10)$$

Here, the correlation function is given by the following integral

$$C(t-s) = \int_0^\infty (2n_{\mathrm{B}}(\omega)+1)J(\omega)\cos[\omega(t-s)]d\omega, \quad (11)$$

where the Bose-Einstein occupation number and the spectral density are given by

$$n_{\mathrm{B}}(\omega) = \left(e^{\frac{\hbar\omega}{k_{\mathrm{B}}T}} - 1\right)^{-1}, \quad J(\omega) = 2\sum_k |f_k|^2\delta(\omega-\omega_k), \quad (12)$$

respectively. Here k_{B} is the Boltzmann constant. In the following we assume that the qubit is initially in an equal superposition state.

The specific form of the coupling constants, and consequently the specific form of the spectral density, depends on the physical problem under study. We follow Ref. [37] and take

$$J(\omega) = g\frac{\omega^p}{\Lambda^{p-1}}\exp\left(-\frac{\omega}{\Lambda}\right), \quad (13)$$

with the cut-off frequency $\Lambda = 1$. p denotes the Ohmicity parameter; for $p \in [0,1)$ we are dealing with a sub-Ohmic environment, for $p = 1$ the environment is Ohmic, and for $p > 1$ it is super-Ohmic. The features that can be displayed by the evolution of the qubit qualitatively depend on the value of this parameter. The parameter g is responsible for the overall strength of the coupling and in this section is taken $g = 1$.

To compute the correlation function, we are tasked with computing the integral over ω exactly. To this end, denote $k(\omega,t) = J(\omega)\cos(\omega t)$ so that we can rewrite the correlation function as

$$\begin{aligned} C(t) &= 2\int_0^\infty k(\omega,t)d\omega + \int_0^\infty k(\omega,t)n_{\mathrm{B}}(\omega)d\omega \\ &:= I_1(t) + I_2(t). \end{aligned}$$

With assumed $\hbar = 1$ and $k_{\mathrm{B}}T = 1$, the first integral is found to be

$$I_1(t) = \Gamma(p+1)\left(t^2+1\right)^{-\frac{p}{2}-\frac{1}{2}}\cos\left((p+1)\tan^{-1}(t)\right),$$

while the second integral can be handled using special functions by observing that

$$\begin{aligned} I_2(t) &= 2\int_0^\infty \frac{e^{-\omega}\omega^p\cos(\omega t)}{e^\omega - 1}d\omega \\ &= 2\,\mathrm{Re}\left(\int_0^\infty \frac{\omega^p e^{-(2-it)\omega}}{1-e^{-\omega}}d\omega\right) \\ &= 2\,\mathrm{Re}\left(\Gamma(p+1)\zeta(p+1,2-it)\right), \end{aligned}$$

where Γ is the standard gamma function, and ζ is the Hurwitz-Riemann zeta function. We see clearly that, for this learning problem, regression on the space of polynomials is likely to be insufficient for reproducing the behavior of these special functions. For this reason, we introduce a Pade approximant, given by Equation (4). Specifically, we use a [4/4] Pade approximant and, consequently, try to learn ten real parameters.

We prepare synthetic data on a time domain of $t \in [10^{-6}, 3]$ with 64 grid points. Using an initial time of 10^{-6} is done to avoid the removable singularity of Equation (11) at $t = 0$. To numerically construct the correlation function (11), we use MATLAB's built-in Riemann-Hurwitz zeta function by evaluating `hurwitzZeta(p+1, 2-1i*t)`, along with evaluating the remaining special and elementary functions in standard ways. We then numerically integrate Equation (10), from an initial condition

of $\rho_{01}(0) = \frac{1}{2}$, by using the non-local Crank-Nicholson scheme outlined in Appendix A.

For this problem, we observe empirically through numerical experimentation that a regularization on the correlation function is unnecessary. Therefore, we solve the unregularized optimization problem (8), that is, with $\alpha = 0$, and for each of the three different physical regimes governed by the parameter Ohmic spectral parameter p. We study the subohmic $0 < p < 1$, ohmic ($p = 1$), and superohmic regimes ($p > 1$), and in all examples, accurately reproduce the dynamics furnished by $C(t)$ using the optimal Pade approximant $P(t, \xi^*)$. Our results are visualized by Figure 1.

Indeed, we observe a typical loss, that is $\mathcal{J}_{\text{reg}}[\xi_*]$, on the order of 10^{-9}. Moreover, the realization of the dynamics, by virtue of the fast evaluations of Pade approximants, is on the order of hundredths of a second on an average Macbook Air laptop. Meanwhile, the synthetic data took about 6 seconds to prepare. Furthermore, the reconstruction of the correlation function in all three parametric regimes is accurate, despite only being trained on one trajectory, therefore, we do not investigate how our trained model generalizes given the uniqueness of integral curves from the dynamics.

Now, we investigate the effect of corrupting our synthetic data with noise. We perturb the Ohmic data, which can be seen in the top panel of Figure 1 with a 10% relative amplitude, that is, $\rho_{01}^{w}(t) = (1 + w(t))\rho_{01}^{\text{data}}(t)$ where $w(t)$ is uniformly sampled on $[-1/10, 1/10]$ at each time t. The trained dynamics, without regularization, typically yields a loss that evaluates on the order of 10^{-4}.

Despite this somewhat small loss, we observe that the correlation function begins to exhibit a fast oscillation near $t = 0.5$. To regularize this behavior, we use a homogeneous Sobolev regularization ($\beta = 1$ in Equation (7)) while varying the weighting parameter α in eq. (8). As can be seen in Figure 2, the regularization does well to smooth out this fast oscillation, yet maintains an evaluation of the loss function that remains within the same order of magnitude as the unregularized case.

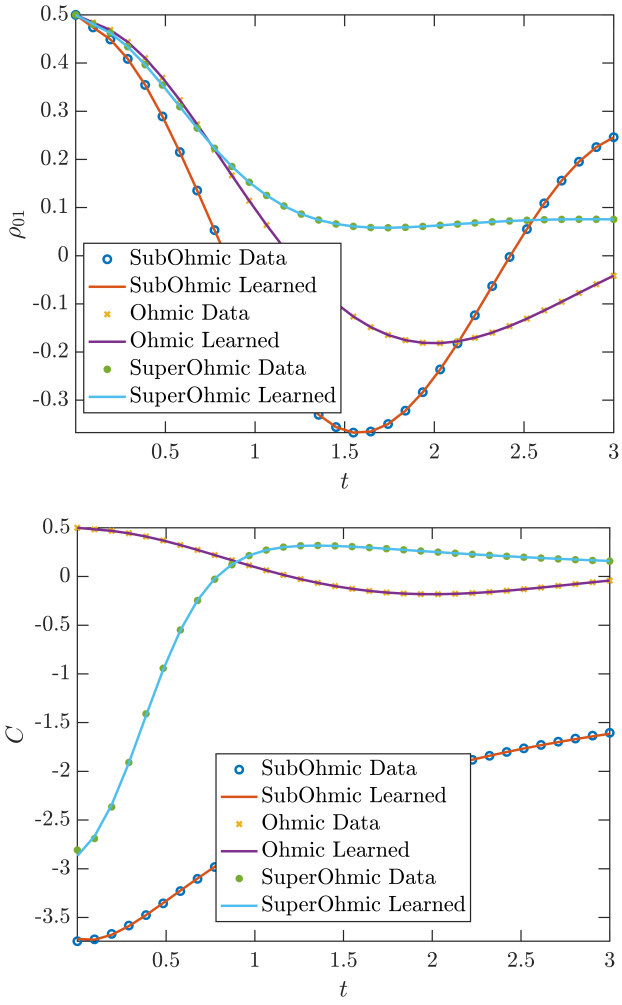

FIG. 1. A numerical solution of Problem (8), with $\alpha = 0$ and constrained by Equation (10), using the [4/4] Pade approximant expressed by Equation (4). We use a subohmic parameter of $p = 1/2$ and a superohmic parameter of $p = 2$. The construction of the synthetic data used in this study is discussed in the main text.

IV. TEST PROBLEM 2: BEYOND THE SPIN-BOSON MODEL: DECAY AND DECOHERENCE

A similar qubit-environment Hamiltonian to that of Test Problem 1 can be used to model an interaction that does involve energy exchange and leads not only to pure dephasing, but involves the evolution of qubit occupations in any basis [38–40]. The condition is that the free qubit Hamiltonian does not commute with the interaction term [41]. To fulfill this condition, we exchange the interaction in eq. (9) as follows,

$$\sigma_z \sum_k \left(f_k b_k^\dagger + f_k^* b_k \right) \to \sigma_x \sum_k \left(f_k b_k^\dagger + f_k^* b_k \right). \quad (14)$$

Note, that the change in the Hamiltonian is small (limited to the exchange of the Pauli matrix that governs the effect of the interaction on the qubit), but it leads to a fundamental change of the nature of the decoherence.

Thus, now both the occupations and the coherences of the qubit state affected by the interaction with the environment, and their evolution (in the interaction picture) is governed by the following equations,

$$\frac{d\rho_{00}}{dt} = -\int_0^t ds \int_0^\infty d\omega\, J(\omega) \cos\left((\varepsilon_0 - \omega)(t-s)\right) \times \left[\rho_{00}(s)(2n(\omega)+1) - \rho_{11}(s)(n(\omega)+1)\right], \quad (15a)$$

$$\frac{d\rho_{01}}{dt} = -\frac{1}{2}\int_0^t ds \int_0^\infty d\omega\, J(\omega)(2n(\omega)+1) \times \left[\rho_{01}(s)e^{-i(\varepsilon_0-\omega)(t-s)} - \rho_{10}(s)e^{-2i\varepsilon_0 t}e^{i(\varepsilon_0-\omega)(t-s)}\right], \quad (15b)$$

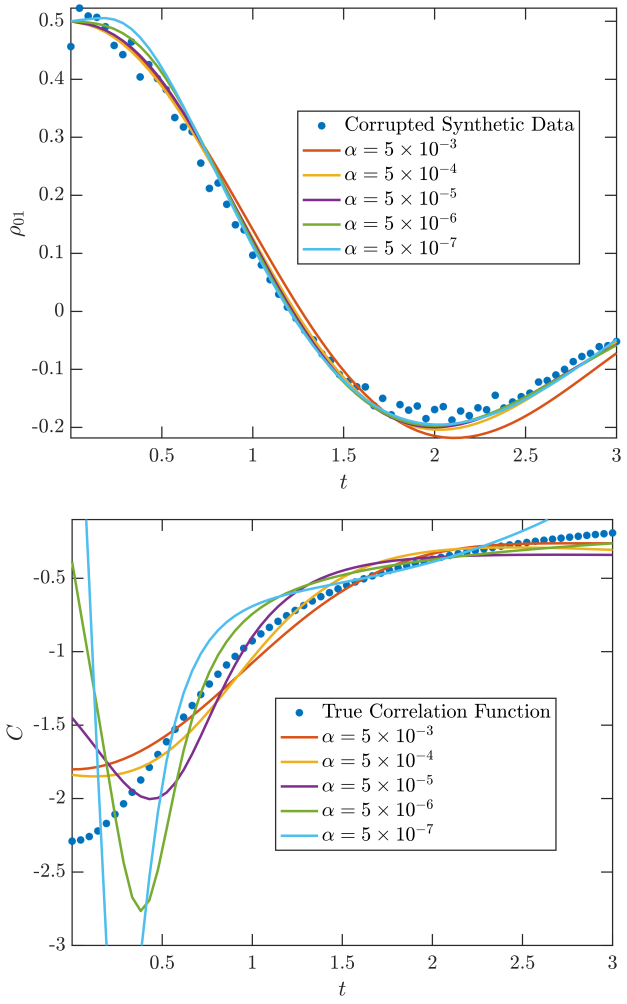

FIG. 2. We display the effect of noise on the regression by performing the same study as in the Ohmic case shown in Figure 1, yet introduce corruption in the synthetic data in the way described in the main text. We observe that even a small Tikhonov regularization does well to smooth out singular features of the correlation function, while maintaining a fairly reasonable fit to the synthetic data. Decreasing the regularization leads to a montonically decreasing singular feature in the bottom panel near $t = 0.5$.

together with the closure relations $\rho_{00}(t) + \rho_{11}(t) = 1$ and $\rho_{10}(t) = \rho_{01}^*(t)$. The form of eqs (15) already holds the assumptions that the initial state of the environment is a thermal equilibrium state with respect to its free Hamiltonian.

We recognize that the dynamics can be written in the form of eq. (3), $A = \mathbf{0}_{4\times4}$ and where B is of the form

$$B = \begin{pmatrix} B_{00} & B_{11} & 0 & 0 \\ -B_{00} & -B_{11} & 0 & 0 \\ 0 & 0 & B_{01} & B_{10} \\ 0 & 0 & B_{10}^* & B_{01}^* \end{pmatrix}. \tag{16}$$

This desirable block matrix structure in the correlation kernel is a major reason why we vectorize the density matrix in the specific way chosen in context of eq. (3).

In the following, we use almost all of the same parameters as in Test Problem 1. Of the parameters that enter the spectral density given by eq. (13), we only change the parameter g to $g = 1/4$ responsible for the overall strength of the coupling, in order to be in the weak coupling limit. We also set the qubit energy splitting $\varepsilon_0 = 1$; this parameter did not affect the evolution in the previous example (while in the interaction picture), thus it was not set. Initial states are drawn from the physical set by sampling

$$R = \begin{pmatrix} s_1 & is_2 \\ -is_2 & 1 - s_1 \end{pmatrix}, \qquad s_1, s_2 \sim \text{Uniform}(0, 1),$$

which is an effective projection onto the positive semidefinite cone with unit trace before vectorizing to the initial state vector $x_0 \in \mathbb{C}^4$.

Generating synthetic data is now more expensive than in the previous case. We must compute quadratures over the angular frequency ω to numerically access the correlation functions. Since these functions must be recomputed at every instance of time t, the full dynamics are costly to construct. To accurately generate the data within an absolute error of 10^{-6} at time $t = 3$ using 32 grid points, we must use a frequency cutoff of $\Omega = 1000$ and 2^{18} uniformly spaced frequency points for a numerical integration over ω using the trapezoidal rule.

To learn each of the four correlation functions, we found it sufficient to use a [3/3] Pade approximant. Empirically for this case study, we found regularization parameters of $\alpha = 10^{-4}$ and $\beta = 0.95$ worked sufficiently well to dampen sharp oscillations almost surely to appear without regularization. We mention in passing that we found it difficult to find a good balance between α and $\beta = 1$, and this is the main reason why we introduce the parameter β into the regularization given by eq. (7).

We also find that the dynamics generated by our learned correlation functions do not generalize from learning a single trajectory. Instead, we train on 30 trajectories evolved from initial conditions sampled from the space ϱ of density matrices, that is, initial conditions that lead to a $\rho(0)$ with unit trace and real, positive eigenvalues. To discuss generalization, we use the empirical risk

$$\mathcal{R}[\tilde{A}, \tilde{B}] = \mathbb{E}_{\rho(0)\sim\varrho}\mathcal{J}_{\text{loss}}[\tilde{A}, \tilde{B}]. \tag{17}$$

We observe an empirical risk evaluated over 100 out of training samples on the order of 10^{-4}, which is the same order of magnitude as our loss function evaluated over the training set. Therefore, we conclude that our trained model generalizes to out of training samples despite an inaccurate reconstruction of the correlation function. We summarize these results visually in Figure 3.

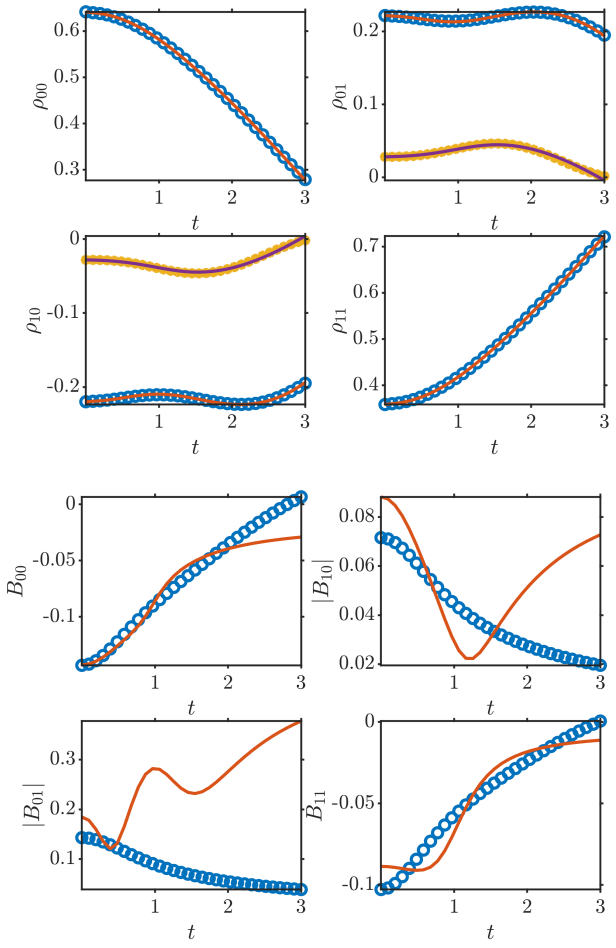

FIG. 3. In the top panel, we display the dynamics, governed by eqs (15), evolved from an out of training sample in the space of initial density matrices ϱ and generated by the learned correlation functions in bottom panel. In the off-diagonal components, real parts correspond to open circles while imaginary parts correspond to solid circles. We observe that despite a reasonable generalization and that Tikhonov regularization ensures smoothness, the learned correlation functions are not close by any metric to the numerically evaluated correlation functions embedded in eqs (15).

V. TEST PROBLEM 3: NONCOMMUTING NOISE AND CROSS–CORRELATED CHANNELS

The final test problem illustrates the full structure of non-Markovian effects when the environment couples through *noncommuting* operators. Unlike the previous examples, where memory kernels act through a single dephasing channel or preserved a block–diagonal structure, simultaneous longitudinal and transverse couplings generate population–coherence transfer, phase–sensitive relaxation, and nontrivial mixing between dynamical modes. Such behavior is typical in solid-state qubits subject to both charge- and spin-type fluctuations, or whenever the bath couples via operators that do not share a common eigenbasis [1, 3, 42].

We work in the interaction picture with respect to $(\varepsilon/2)\sigma_z$, so that coherent precession appears inside the kernel through the rotating operator

$$\sigma_x(t) = \cos(\Delta t)\sigma_x + \sin(\Delta t)\sigma_y.$$

Under the standard Born approximation [1, 4, 5], the Nakajima–Zwanzig generator for simultaneous σ_x– and σ_z–couplings has the double–commutator structure

$$\begin{aligned}\mathcal{K}(t)\rho = {} & C_{xx}(t)\,[\sigma_x,[\sigma_x(t),\rho]] \\ & + C_{zz}(t)\,[\sigma_z,[\sigma_z,\rho]] \\ & + C_{xz}(t)\,[\sigma_x,[\sigma_z,\rho]] \\ & + \overline{C_{xz}(t)}\,[\sigma_z,[\sigma_x(t),\rho]],\end{aligned} \tag{18}$$

where the scalar kernels $C_{\alpha\beta}(t)$ encode the bath correlations and satisfy $C_{zx}(t) = C_{xz}(t)^*$ by Hermiticity.

Vectorizing the density matrix as in Sec. II,

$$x(t) = \big[\rho_{00}(t), \rho_{11}(t), \rho_{01}(t), \rho_{10}(t)\big]^\top,$$

the reduced dynamics take the Volterra form

$$\dot{x}(t) = \int_0^t B(t,\tau)\,x(\tau)\,d\tau, \tag{19}$$

with $A = 0$ and

$$\begin{aligned}B(t,\tau) = {} & C_{xx}(t-\tau)\,B_{xx}(t) + C_{zz}(t-\tau)\,B_{zz}(t) \\ & + C_{xz}(t-\tau)\,B_{xz}(t) + \overline{C_{xz}(t-\tau)}\,B_{zx}(t),\end{aligned} \tag{20}$$

where the 4×4 matrices $B_{\alpha\beta}(t)$ follow directly from Eq. (18). This is the first example in which the kernel is fully non-block-diagonal and mixes populations with coherences.

To stress-test the learning procedure, we prescribe synthetic kernels with distinct physical time scales,

$$\begin{aligned}C_{xx}(t) &= \alpha_x e^{-\gamma_x t}\cos(\Omega_x t),\\ C_{zz}(t) &= \alpha_z e^{-\gamma_z t},\\ C_{xz}(t) &= \kappa\, e^{-\gamma_c t}\, e^{-i\phi},\end{aligned}$$

producing oscillatory memory in the transverse channel, simple decay in the longitudinal channel, and a phase–sensitive cross-term. For numerical purposes, we introduce a short exponential ramp to regularize the removable singularity that appears as $t \downarrow 0$.

We integrate Eq. (19) on $t \in [10^{-6}, 3]$ using $M = 32$ grid points and the nonlocal Crank–Nicolson scheme described in Appendix A. To enforce complete positivity of the scalar kernels during learning, we adopt a rank–1 factorization

$$\begin{aligned}C_{xx}(t) &= f_x(t)\,f_x(t)^*,\\ C_{zz}(t) &= f_z(t)\,f_z(t)^*,\\ C_{xz}(t) &= f_x(t)\,f_z(t)^*.\end{aligned}$$

so that all channels share a common latent representation. Each f_α is modeled by a [3/3] Padé approximant

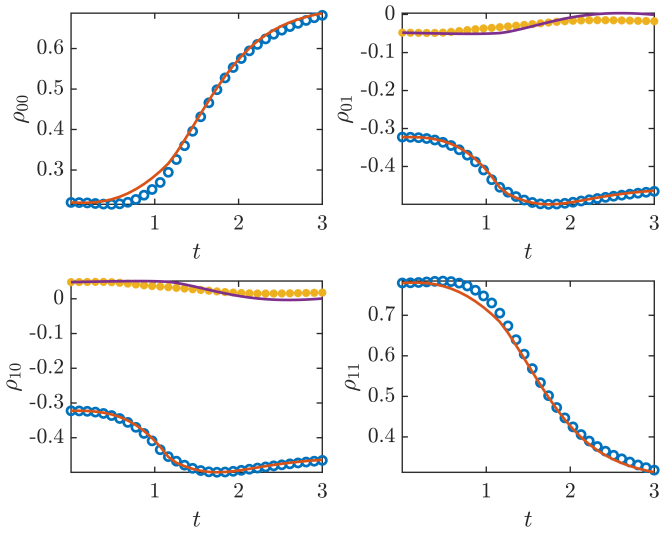

FIG. 4. Just as in Figure 3, we display the dynamics, evolved from an out of training sample in the space of initial density matrices ϱ and generated by learned correlation functions. This shows that a rational hypothesis of the correlation kernel is able to reproduce the dynamics of the density matrix on moderate time scales.

as in Eq. (4), which provides a compact parameterization while allowing complex oscillatory and decaying behavior. The full kernel is then assembled through Eq. (20).

The objective functional follows the same structure as Eq. (8): an L^2 misfit between learned and synthetic trajectories, averaged over the training set, together with a Tikhonov term on the scalar kernels. As in Test Problem 2, only an extremely small regularization weight is required to suppress spurious high-frequency structure. Optimization is performed once again using `fminunc` with quasi-Newton updates. Figure 4 shows that the learned dynamics accurately reproduce the held-out trajectory in cross-validating across all components for short times, including the rotating-frame modulation and the cross-channel phase. The cross-validation loss here is on the order of 10^{-3}.

Pointwise kernel recovery is, as in problem 2, significantly more challenging. Many distinct kernels generate nearly identical state trajectories over a finite observation window. Nevertheless, the rank–1 Padé parameterization captures the qualitative temporal structure.

VI. CONCLUSIONS

We proposed a direct, data driven route to non Markovian dynamics by learning the Volterra kernel in the Nakajima Zwanzig representation. The guiding principle was to keep the hypothesis class minimal: each scalar component of the kernel is modeled by a fixed-order Padé rational function in the lag variable, and the search is regularized only by an H^1 penalty. In this scalar setting the learning problem is straightforward. Fits are stable, trajectories are accurate, and the optimization landscape behaves predictably. The difficulty only appears once these well behaved scalar components are assembled into a matrix-valued kernel. At that point the familiar non identifiability issues arise, and the behavior can no longer be attributed to the Padé model itself but rather to the structure induced by the operator-valued coupling.

Within this simple framework, the nonlocal Crank Nicolson discretization already recovers accurate state trajectories across three increasingly challenging test beds: pure dephasing, a generalized spin boson model with decoherence and population decay, and fully cross correlated, noncommuting channels. On the analysis side, we established regularity for $B \in H^1$ and existence of minimizers for the regularized problem, which gives a basic well posedness foundation for learning operator-valued kernels.

The approach has clear limitations. First, identifiability of (A, B) from finite time windows is delicate. Distinct operator-valued kernels can generate nearly indistinguishable trajectories on $[0, T]$, and our results show that accurate state fits do not guarantee pointwise kernel recovery. Second, we did not enforce complete positivity. Trace preservation and Hermiticity are respected by construction, but CP is only implicit and may fail outside the training window. Third, Padé models can introduce spurious poles and local oscillations near the origin in noisy settings. The H^1 penalty mitigates this behavior but does not remove it completely. Finally, the nonlocal stepping with a dense history is computationally heavy, with quadratic scaling in N_t, which limits long horizon training and broad hyperparameter sweeps.

These observations suggest several concrete extensions. On the kernel side, barycentric rational representations [43, 44] (or vector–fitting style constraints) could replace Padé to improve pole placement, enable degree adaptivity, and stabilize the small–lag regime. On the structure side, one can enforce GKSL–compatible parametrizations at the memory level [24, 25, 42], e.g., by learning positive semidefinite Kossakowski tensors in a double–commutator basis or by incorporating transfer–tensor constraints [15].

On the data and experimental side, longer training windows and multi–trajectory ensembles (temperatures, couplings) are suggested to help pin down kernel tails and reduce nonuniqueness [1–3]. To address this more expensive computation, one may consider fast history evaluation via convolution–quadrature/FFT methods [45] and sum–of–exponentials compressions for kernels to reduce cost from $O(N_t^2)$ to near $O(N_t \log N_t)$, making larger grids and tighter tolerances practical.

With all of this in mind, we invite others to take this minimal baseline as a starting point, preserve the simple-first ethos, and advance structure preserving, scalable, and provably stable methods for interpretable learning of non-Markovian dynamics in open quantum settings.

VII. ACKNOWLEDGMENTS

J.A. acknowledges support from NSF award number 2316622. The contribution of K.R. was supported within the QuantERA II Programme that has received funding from the EU H2020 research and innovation programme under Grant Agreement No 101017733, and with funding organisation MEYS (The Ministry of Education, Youth and Sports) of the Czech Republic.

Appendix A: Numerical Method for Solving Volterra Integro-Differential Equations

We consider the numerical solution of the Volterra-type integro–differential equation

$$\frac{d\mathbf{x}}{dt} \;=\; A\mathbf{x}(t) \;+\; \int_{t_0}^{t} B(t-\tau)\,\mathbf{x}(\tau)\,d\tau \;=:\; \mathbf{f}(\mathbf{x},t). \quad \text{(A1)}$$

This formulation permits the application of standard quadrature methods. Indeed, using the trapezoidal rule, one obtains the Crank–Nicolson update

$$\begin{aligned}\mathbf{x}_{n+1} &= \mathbf{x}_n + \int_{t_n}^{t_{n+1}} \mathbf{f}(\mathbf{x},\tau)\,d\tau \\ &= \mathbf{x}_n + \tfrac{h}{2}\big(\mathbf{f}_n + \mathbf{f}_{n+1}\big) + \mathcal{O}(h^3), \qquad \text{(A2)}\end{aligned}$$

where the time grid is uniform with step size $h := t_{n+1} - t_n$, $\mathbf{x}(t_n) = \mathbf{x}_n$, and $\mathbf{f}(\mathbf{x}_n, t_n) =: \mathbf{f}_n$. Expanding the right-hand side gives

$$\begin{aligned}\mathbf{f}_n + \mathbf{f}_{n+1} &= A\mathbf{x}_n + A\mathbf{x}_{n+1} \\ &\quad + \int_{t_0}^{t_n} B(t_n - s)\,\mathbf{x}(s)\,ds \\ &\quad + \int_{t_0}^{t_{n+1}} B(t_{n+1} - s)\,\mathbf{x}(s)\,ds. \qquad \text{(A3)}\end{aligned}$$

Approximating these integrals once again by the trapezoidal rule and solving explicitly for $\mathbf{x}_{n+1}$ leads to the nonlocal Crank–Nicolson scheme

$$\mathbf{x}_{n+1} = \Big(I - \tfrac{h}{2}A - \tfrac{h^2}{4}B(t_0)\Big)^{-1}\Big(\mathbf{x}_n + \tfrac{h}{2}\,\mathbf{g}_n\Big), \quad \text{(A4)}$$

where the nonlocal contribution is

$$\begin{aligned}\mathbf{g}_n = A\mathbf{x}_n + \tfrac{h}{2}\Big(B(t_n)\mathbf{x}_0 + 2\sum_{k=1}^{n-1} B(t_n - t_k)\,\mathbf{x}_k + B(t_0)\mathbf{x}_n \\ + B(t_{n+1})\mathbf{x}_0 + 2\sum_{k=1}^{n} B(t_{n+1} - t_k)\,\mathbf{x}_k\Big). \qquad \text{(A5)}\end{aligned}$$

This update formula makes explicit the history dependence of the scheme: each new state $\mathbf{x}_{n+1}$ depends not only on $\mathbf{x}_n$ but on the entire trajectory $\{\mathbf{x}_k\}_{k=0}^{n}$ through weighted contributions of the kernel $B(\cdot)$.

Appendix B: Well-posedness of the Optimization Problem

In this section we establish a theorem that states the optimization problem used to learn the operators A and B is mathematically well-posed. We show that for admissible kernels B the state equation admits a unique solution with sufficient regularity, and that the learning functional $\mathcal{J}$ admits at least one minimizer. The argument follows the direct method in the calculus of variations

We first establish regularity and an a priori estimate for the Volterra state equation; this will control the states associated with a minimizing sequence for the learning functional. Throughout, $\|\cdot\|_2$ denotes the Euclidean norm on $\mathbb{C}^4$, and for matrices $A \in \mathbb{C}^{4\times 4}$ we write $\|A\|$ for the spectral (matrix 2-) norm, i.e., the largest singular value of A.

Lemma: Regularity and a priori estimate. *Let $(A, B) \in \mathcal{O}$ and consider*

$$\dot{\mathbf{x}}(t) = A\,\mathbf{x}(t) + \int_0^t B(t-\tau)\,\mathbf{x}(\tau)\,d\tau, \qquad \mathbf{x}(0) = \mathbf{x}_0 \in \mathbb{C}^4. \quad \text{(B1)}$$

If $B \in H^1\big(0,T;\mathbb{C}^{4\times 4}\big)$, then $\mathbf{x} \in H^1\big(0,T;\mathbb{C}^4\big)$ and

$$\|\mathbf{x}\|_{H^1(0,T)} \;\le\; C\big(\|A\|,\, \|B\|_{H^1(0,T)},\, T,\, \|\mathbf{x}_0\|_2\big)\,. \quad \text{(B2)}$$

Proof. First note that since $H^1(0,T) \hookrightarrow C([0,T])$, there exists a $c_T > 0$ such that $\|B\|_{L^\infty(0,T)} \le c_T \|B\|_{H^1(0,T)}$. Now, for convenience, we define the Volterra operator

$$(\mathcal{K}_B\mathbf{y})(t) := \int_0^t B(t-\tau)\,\mathbf{y}(\tau)\,d\tau. \quad \text{(B3)}$$

We seek to apply Young's inequality for convolutions, which is typically stated for scalar convolutions. Since $B(\cdot)$ is matrix–valued and $x(\cdot)$ is vector–valued, we estimate pointwise using the spectral (matrix 2-) norm of B and the Euclidean norm on $\mathbb{C}^4$:

$$\begin{aligned}\|(\mathcal{K}_B x)(t)\|_2 &= \left\| \int_0^t B(t-\tau)\,x(\tau)\,d\tau \right\|_2 \\ &\le \int_0^t \|B(t-\tau)\|\,\|x(\tau)\|_2\,d\tau. \qquad \text{(B4)}\end{aligned}$$

Define the scalar functions $b(s) := \|B(s)\|$ and $g(s) := \|x(s)\|_2$. Then

$$\|(\mathcal{K}_B x)(\cdot)\|_2 \;\le\; b * g \quad \text{on } [0,T]. \quad \text{(B5)}$$

Thus by Young's inequality on $[0,T]$ with $(p,q,r) = (1,2,2)$,

$$\begin{aligned}\|\mathcal{K}_B x\|_{L^2(0,T)} &\le \|b * g\|_{L^2(0,T)} \\ &\le \|b\|_{L^1(0,T)}\,\|g\|_{L^2(0,T)}\end{aligned}$$

$$= \left(\int_0^T \|B(s)\|\, ds \right) \|x\|_{L^2(0,T)}$$
$$= \|B\|_{L^1(0,T)} \|x\|_{L^2(0,T)}. \tag{B6}$$

Since $H^1(0,T) \hookrightarrow C([0,T])$ and from Hölder's inequality, it follows that

$$\|B\|_{L^1(0,T)} \le \sqrt{T}\, \|B\|_{L^2(0,T)} \le c_T \|B\|_{H^1(0,T)}. \tag{B7}$$

Combining (B6) and (B7) yields

$$\|\mathcal{K}_B x\|_{L^2(0,T)} \le c_T \|B\|_{H^1(0,T)} \|x\|_{L^2(0,T)}. \tag{B8}$$

To estimate norms of the state vector, we use the integral formulation of the state equation (B1) which is

$$\mathbf{x}(t) = \mathbf{x}_0 + \int_0^t \Big(A\,\mathbf{x}(s) + (\mathcal{K}_B \mathbf{x})(s) \Big)\, ds. \tag{B9}$$

Taking Euclidean norms and applying (B8), for $t \in [0,T]$ we obtain

$$\|\mathbf{x}(t)\|_2 \le \|\mathbf{x}_0\|_2 + \int_0^t \|A\|\, \|\mathbf{x}(s)\|_2\, ds + \int_0^t \|(\mathcal{K}_B \mathbf{x})(s)\|_2\, ds$$
$$\le \|\mathbf{x}_0\|_2 + \int_0^t \Big(\|A\| + c_T \|B\|_{H^1(0,T)} \Big) \|\mathbf{x}(s)\|_2\, ds. \tag{B10}$$

Finally, applying Grönwall's inequality to (B10) gives the uniform bound

$$\|\mathbf{x}\|_{L^\infty(0,T;\mathbb{C}^4)} \le \|\mathbf{x}_0\|_2 \exp\Big(T\big(\|A\| + c_T \|B\|_{H^1(0,T)} \big) \Big). \tag{B11}$$

□

Existence of minimizers and sequential continuity. *Let $\mathcal{O}$ be as in (5), with $B \in H^1(0,T;\mathbb{C}^{4\times 4})$ and A ranging over a nonempty, closed, and bounded subset of $\mathbb{C}^{4\times 4}$. Given data $\mathbf{x}^{\mathrm{data}} \in L^2(0,T;\mathbb{C}^4)$ and an initial condition $\mathbf{x}(0) = \mathbf{x}_0 \in \mathbb{C}^4$, the constrained optimization problem (8) admits at least one minimizer $(A_*, B_*) \in \mathcal{O}$.*

Moreover, if $(A_k, B_k) \rightharpoonup (A_*, B_*)$ in $\mathcal{O}$ with $\{(A_k, B_k)\}$ bounded, then the associated states $\mathbf{x}_k$ converge (up to subsequences) strongly in $L^2(0,T)$ to the state $\mathbf{x}_*$ corresponding to (A_*, B_*). In particular, the loss functional $\mathcal{J}_{\mathrm{loss}}$ is sequentially continuous along bounded sequences in $\mathcal{O}$.

Proof. Our first step is to show coercivity of the regularizer. Recall that the regularization term takes the form

$$\mathcal{J}_{\mathrm{reg}}[B] \;=\; (1-\beta)\, \|B\|^2_{L^2(0,T)} \;+\; \beta\, \|\dot{B}\|^2_{L^2(0,T)}.$$

Coercivity follows directly since

$$\min\{1-\beta, \beta\}\, \|B\|^2_{H^1(0,T)} \;\le\; \mathcal{J}_{\mathrm{reg}}[B] \tag{B12}$$
$$\le\; \max\{1-\beta, \beta\}\, \|B\|^2_{H^1(0,T)}, \tag{B13}$$

implies that the sequence $\{B_k\}$ is bounded in $H^1(0,T)$, and by assumption, $\{A_k\}$ is bounded in $\mathbb{C}^{4\times 4}$.

Now, we aim to show uniform bounds for the states and kernels. To this end, let $\mathbf{x}_k$ be the state associated with (A_k, B_k). Lemma B gives the a priori estimate

$$\|\mathbf{x}_k\|_{H^1(0,T)} \;\le\; C\big(\sup_k \|A_k\|, \sup_k \|B_k\|_{H^1}, T, \|\mathbf{x}_0\| \big).$$

Hence $\{\mathbf{x}_k\}$ is bounded in $H^1(0,T;\mathbb{C}^4)$. By compactness of the embedding $H^1(0,T) \hookrightarrow L^2(0,T)$, a subsequence converges strongly:

$$\mathbf{x}_k \;\to\; \mathbf{x}_* \quad \text{in } L^2(0,T).$$

On the side of the correlation functions, since $\mathbb{C}^{4\times 4}$ is finite-dimensional, bounded sequences admit convergent subsequences: $A_k \to A_*$. For B_k, boundedness in $H^1(0,T)$ implies relative compactness in $C([0,T])$, so $B_k \to B_*$ uniformly on $[0,T]$.

Now, we involve the dynamics. Recall that each state satisfies the integral equation

$$\mathbf{x}_k(t) = \mathbf{x}_0 + \int_0^t \big(A_k \mathbf{x}_k(s) + (\mathcal{K}_{B_k} \mathbf{x}_k)(s) \big)\, ds.$$

For the drift term $A_k \mathbf{x}_k$, convergence follows from $A_k \to A_*$ and $\mathbf{x}_k \to \mathbf{x}_*$ in L^2. For the Volterra term, decompose

$$\mathcal{K}_{B_k} \mathbf{x}_k - \mathcal{K}_{B_*} \mathbf{x}_* = \mathcal{K}_{B_k}(\mathbf{x}_k - \mathbf{x}_*) + (\mathcal{K}_{B_k} - \mathcal{K}_{B_*}) \mathbf{x}_*.$$

The first term vanishes by the uniform bound $\|\mathcal{K}_{B_k}\|_{L^2 \to L^2} \le c\, \|B_k\|_{H^1}$ and strong convergence of $\mathbf{x}_k$. The second vanishes since $B_k \to B_*$ uniformly, which yields $\|B_k - B_*\|_{L^1} \to 0$ and hence operator convergence $\mathcal{K}_{B_k} \to \mathcal{K}_{B_*}$ in $L^2 \to L^2$. Thus the limit $\mathbf{x}_*$ satisfies the state equation with coefficients (A_*, B_*).

Having convergence in hand and to complete the proof, it remains to show lower semicontinuity of the functional. But this is simple to show, by design, since the regularizer $\mathcal{J}_{\mathrm{reg}}$ is convex and weakly lower semicontinuous on $H^1(0,T)$, and the misfit

$$\mathcal{J}_{\mathrm{loss}}[A,B] = \int_0^T \|\mathbf{x}^{\mathrm{data}}(t) - \mathbf{x}(t;A,B)\|^2\, dt$$

is sequentially continuous along bounded sequences thanks to strong L^2 convergence of states. Therefore,

$$\mathcal{J}[A_*, B_*] \;\le\; \liminf_{k\to\infty} \mathcal{J}[A_k, B_k] = \inf_{(A,B)\in\mathcal{O}} \mathcal{J}[A,B],$$

showing (A_*, B_*) is a minimizer. The sequential continuity claim follows from the same convergence argument. □

Remark: The proof above relies on classical ingredients, namely, coercivity of the H^1–regularizer, compactness of Sobolev embeddings, Grönwall-type a priori estimates, and weak lower semicontinuity of convex functionals. These are standard tools in the analysis of Volterra equations and variational problems [46–48].

Nonetheless, establishing existence of minimizers and sequential continuity is essential in our setting. The learning problem is formulated over operator-valued kernels where only weak convergence of approximants is natural (e.g. when passing from Padé rational fits to their limits). Without an a priori guarantee, one could not exclude pathological minimizing sequences or loss of stability under weak convergence. Theorem B ensures that every minimizing sequence admits a subsequence converging to a genuine solution, and that the data misfit functional behaves continuously along bounded operator families. Despite its lack of mathematical novelty, this theoretical foundation justifies the numerical sections of this work and is presented here for sake of completeness.

REFERENCES

[1] H.-P. Breuer and F. Petruccione, *The Theory of Open Quantum Systems* (Oxford University Press, 2002).

[2] I. de Vega and D. Alonso, Dynamics of non-markovian open quantum systems, Rev. Mod. Phys. **89**, 015001 (2017).

[3] A. Rivas and S. F. Huelga, *Open Quantum Systems: An Introduction* (Springer, 2012).

[4] S. Nakajima, On quantum theory of transport phenomena, Prog. Theor. Phys. **20**, 948 (1958).

[5] R. Zwanzig, Ensemble method in the theory of irreversibility, J. Chem. Phys. **33**, 1338 (1960).

[6] H. Mori, Transport, collective motion, and brownian motion, Prog. Theor. Phys. **33**, 423 (1965).

[7] G. Ciccotti, M. H. Kalos, and J. C. Tully, Projection methods in molecular dynamics, J. Stat. Phys. **118**, 373 (2005).

[8] R. Zwanzig, *Nonequilibrium Statistical Mechanics* (Oxford University Press, 2001).

[9] A. J. Chorin and O. H. Hald, *Stochastic Tools in Mathematics and Science*, 3rd ed. (Springer, 2013).

[10] Z. Li, X. Bian, T. A. Caswell, and G. E. Karniadakis, Data-driven parametrization of memory kernels in generalized langevin dynamics, J. Chem. Phys. **146**, 014104 (2017).

[11] G. Jung, M. Hanke, and F. Schmid, Coarse-grained molecular dynamics with memory kernels, J. Chem. Theory Comput. **13**, 2481 (2017).

[12] J. Wang, S. Olsson, C. Wehmeyer, and F. Noé, Machine learning of coarse-grained molecular dynamics force fields, J. Chem. Phys. **152**, 194106 (2020).

[13] F. Noé and C. Clementi, Collective variables for machine learning in molecular dynamics, Nature Reviews Physics **2**, 42 (2020).

[14] F. A. Pollock, C. Rodríguez-Rosario, T. Frauenheim, M. Paternostro, and K. Modi, Operational markov condition for quantum processes, Phys. Rev. Lett. **120**, 040405 (2018).

[15] J. Cerrillo and J. Cao, Non-markovian dynamical maps: Numerical processing of open quantum trajectories, Phys. Rev. Lett. **112**, 110401 (2014).

[16] S. Krastanov and L. Jiang, Stochastic modeling of quantum dynamics with process tensors, Phys. Rev. A **103**, 032605 (2021).

[17] P. Sentz, S. Nicholson, Y. Cho, S. Reddy, B. Keith, and S. Günther, Learning thermodynamic master equations for open quantum systems, arXiv preprint arXiv:2506.01882 (2025), submitted; includes thermodynamically consistent machine learning of open quantum dynamics.

[18] A. J. Leggett, S. Chakravarty, A. T. Dorsey, M. P. A. Fisher, A. Garg, and W. Zwerger, Dynamics of the dissipative two-state system, Rev. Mod. Phys. **59**, 1 (1987).

[19] U. Weiss, *Quantum Dissipative Systems*, 2nd ed. (World Scientific, 1999).

[20] B. Krummheuer, V. M. Axt, and T. Kuhn, Theory of pure dephasing and the resulting absorption line shape in semiconductor quantum dots, Phys. Rev. B **65**, 195313 (2002).

[21] A. J. Ramsay, A. V. Gopal, E. M. Gauger, A. Nazir, B. W. Lovett, A. M. Fox, and M. S. Skolnick, Phonon-induced rabi-frequency renormalization of optically driven single ingaas/gaas quantum dots, Phys. Rev. Lett. **104**, 017402 (2010).

[22] G. Clos and H.-P. Breuer, Quantification of memory effects in the spin-boson model, Phys. Rev. A **86**, 012115 (2012).

[23] C. Flindt, T. Novotný, and A.-P. Jauho, Full counting statistics of nano-electromechanical systems, Europhys. Lett. **69**, 475 (2005).

[24] V. Gorini, A. Kossakowski, and E. C. G. Sudarshan, Completely positive dynamical semigroups of n-level systems, Journal of Mathematical Physics **17**, 821 (1976).

[25] G. Lindblad, On the generators of quantum dynamical semigroups, Communications in Mathematical Physics **48**, 119 (1976).

[26] F. Shibata, Y. Takahashi, and N. Hashitsume, Generalized projection operator formalism for non-equilibrium statistical mechanics, Journal of Statistical Physics **17**, 171 (1977).

[27] H.-P. Breuer, B. Kappler, and F. Petruccione, Stochastic wave-function method for non-markovian quantum master equations, Phys. Rev. A **59**, 1633 (1999).

[28] H.-P. Breuer, B. Kappler, and F. Petruccione, The time-convolutionless projection operator technique in the quantum theory of dissipation and decoherence, Annals of Physics **291**, 36 (2001).

[29] H.-P. Breuer, J. Gemmer, and M. Michel, Non-markovian quantum dynamics: Correlated projection superoperators and hilbert space averaging, Phys. Rev. E **73**, 016139 (2006).

[30] W. H. Zurek, Decoherence, einselection, and the quantum origins of the classical, Rev. Mod. Phys. **75**, 715 (2003).

[31] K. Roszak and L. Cywiński, Characterization and measurement of qubit-environment-entanglement generation during pure dephasing, Phys. Rev. A **92**, 032310 (2015).

[32] K. Roszak, Criteria for system-environment entanglement generation for systems of any size in pure-dephasing evolutions, Phys. Rev. A **98**, 052344 (2018).

[33] A. J. Leggett, S. Chakravarty, A. T. Dorsey, M. P. A. Fisher, A. Garg, and W. Zwerger, Dynamics of the dissipative two-state system, Rev. Mod. Phys. **59**, 1 (1987).

[34] C. Guo, A. Weichselbaum, J. von Delft, and M. Vojta, Critical and strong-coupling phases in one- and two-bath spin-boson models, Phys. Rev. Lett. **108**, 160401 (2012).

[35] Z. Cai, U. Schollwöck, and L. Pollet, Identifying a bath-induced bose liquid in interacting spin-boson models, Phys. Rev. Lett. **113**, 260403 (2014).

[36] L. Ferialdi, Exact non-markovian master equation for the spin-boson and jaynes-cummings models, Phys. Rev. A **95**, 020101 (2017).

[37] A. Lampo, J. Tuziemski, M. Lewenstein, and J. K. Korbicz, Objectivity in the non-markovian spin-boson model, Phys. Rev. A **96**, 012120 (2017).

[38] D. P. DiVincenzo and D. Loss, Rigorous born approximation and beyond for the spin-boson model, Phys. Rev. B **71**, 035318 (2005).

[39] W. Wu and M. Liu, Effects of counter-rotating-wave terms on the non-markovianity in quantum open systems, Phys. Rev. A **96**, 032125 (2017).

[40] B. Gulácsi and G. Burkard, Signatures of non-markovianity of a superconducting qubit, Phys. Rev. B **107**, 174511 (2023).

[41] M. Strzałka, R. Filip, and K. Roszak, Qubit-environment entanglement in time-dependent pure dephasing, Phys. Rev. A **109**, 032412 (2024).

[42] B. Vacchini, A. Smirne, E.-M. Laine, J. Piilo, and H.-P. Breuer, Generalized master equations for non-markovian dynamics, New Journal of Physics **13**, 093004 (2011).

[43] J.-P. Berrut and L. N. Trefethen, Barycentric lagrange interpolation, SIAM Rev. **46**, 501 (2004).

[44] Y. Nakatsukasa, O. Sète, and L. N. Trefethen, The aaa algorithm for rational approximation, SIAM J. Sci. Comput. **40**, A1494 (2018).

[45] C. Lubich, Convolution quadrature and discretized operational calculus. i, Numer. Math. **52**, 129 (1988).

[46] H. Brezis, *Functional Analysis, Sobolev Spaces and Partial Differential Equations* (Springer, 2010).

[47] E. Zeidler, *Nonlinear Functional Analysis and its Applications II: Variational Methods and Optimization* (Springer, 1990).

[48] L. C. Evans, *Partial Differential Equations*, 2nd ed. (American Mathematical Society, 2010).