

Learning Arbitrary Lindbladians with Quantum Error Correction

Nikita Romanov,^{1,*} Petr Ivashkov,^{1,2} Weiyuan Gong,³
Ishaan Kannan,¹ Andi Gu,¹ Hong-Ye Hu,¹ and Susanne F. Yelin^{1,†}

¹*Department of Physics, Harvard University, Cambridge, MA 02138, USA*

²*Department of Information Technology and Electrical Engineering, ETH Zürich, Zürich, Switzerland*

³*School of Engineering and Applied Sciences, Harvard University, Allston, MA 02134, USA*

In this work, we advance the characterization of open quantum systems by giving, to the best of our knowledge, the first protocols for learning sparse Lindbladians that attain the Heisenberg limit for coherent components and the standard quantum limit for the full Lindbladian, all without prior knowledge of the underlying interaction structure.

Motivation. Reconstructing the description of a quantum system from observations of its dynamics is a foundational task across quantum science, including benchmarking of noisy devices, metrology in the presence of decoherence and robust quantum control. Under Markovian noise, quantum dynamics are generated by a Lindbladian, which covers both coherent Hamiltonian evolution and dissipation [1, 2]. This gives rise to two complementary learning tasks. Hamiltonian learning targets the coherent part of the generator, treating dissipation as noise that can obscure or bias estimates of the Hamiltonian parameters. Lindbladian learning is the strictly harder task of reconstructing the full generator, including its dissipative components. In this work, we consider both of these learning tasks in the *ansatz-free* setting, where the relevant Hamiltonian and dissipator terms are not specified in advance. This setting is crucial for realistic hardware, where coherent control imperfections and dissipative error mechanisms may be unknown and nonlocal, leaving the learner to contend with an exponentially large candidate space.

Gaps in the literature. In particular, we highlight two key open problems. The first one is attaining the Heisenberg-limited (HL) evolution time for *ansatz-free* Hamiltonian learning in the presence of decoherence. Achieving the HL is a long-standing goal in quantum sensing and learning, but this scaling is notoriously fragile under noise [3, 4]. This fragility inspired extensive work on quantum error-corrected metrology, where codes are tailored to a pre-specified noise model and target Hamiltonian, thereby recovering HL scaling for Hamiltonian parameter estimation [5–7]. This strategy, however, does not readily extend to the *ansatz-free* setting as it is unclear which components of the dynamics should be suppressed and which should be preserved. This motivates a fundamental question addressed by our work:

Q1 : Can one attain *Heisenberg-limited* noisy Hamiltonian learning in the *ansatz-free* setting?

In practice, however, characterizing an open quantum system often requires learning the full Lindbladian, including both coherent and dissipative components. For this task, HL scaling is known to be unachievable, so we target the best possible scaling in this setting—the quadratically worse standard quantum limit (SQL). Despite recent progress on multiqubit Lindbladian learning, SQL-scaling protocols remain limited to local models with known structure and single-qubit dissipation [8, 9], while existing *ansatz-free* approaches do not yet achieve SQL scaling [10]. This motivates the second open question we address:

Q2 : Can one attain *standard-quantum-limited* Lindbladian learning in the *ansatz-free* setting?

Main advance. We answer both questions affirmatively by designing efficient end-to-end algorithms for learning arbitrary Lindbladians with polynomially many terms. Our main technical innovation is a recursive QEC-based framework that adaptively constructs families of stabilizer codes to suppress the strongest identified Lindbladian terms while preserving access to weaker unknown ones. This framework is what enables the HL scaling for *ansatz-free* learning of the Hamiltonian component disjoint from the dissipator and the SQL scaling for learning the full Lindbladian. We also introduce a coefficient-learning primitive that avoids global linear-system inversion, preventing potentially exponential sample-complexity overheads for highly nonlocal Lindbladian terms.

Model and goal. We consider an unknown time-independent n -qubit Lindbladian $\mathcal{L}$ expressed in the Pauli basis as:

$$\mathcal{L}(\rho) = -i \sum_{P_i \in \mathcal{S}_H} h_i [P_i, \rho] + \sum_{P_i, P_j \in \mathcal{S}_D} a_{ij} (P_i \rho P_j - \tfrac{1}{2} \{P_j P_i, \rho\}), \quad (1)$$

* nromanov@g.harvard.edu

† syelin@g.harvard.edu




FIG. 1. **Overview of the learning framework.** Given black-box access to $e^{t\mathcal{L}}$ and a sparsity upper bound M , the algorithm first identifies the heavy dissipator footprint down to a threshold $\eta(\varepsilon)$. The choice of threshold depends on the target task: $\eta = \mathcal{O}(\varepsilon/M)$ suffices for Heisenberg-limited HDD Hamiltonian learning under the balanced-tail condition, while $\eta = \mathcal{O}(\varepsilon^2/M)$ enables full Lindbladian learning at the standard quantum limit. The key primitive is recursive QEC reshaping: previously identified dissipative terms are made correctable, previously identified Hamiltonian terms are made detectable, and random stabilizer codes preserve sensitivity to weaker unknown terms.

where $\mathcal{S}_H$ and $\mathcal{S}_D$ are unknown Pauli supports of the Hamiltonian and dissipator. We denote by M an upper bound on the number of nonzero Pauli-basis coefficients. This upper bound is used only to set algorithmic parameters, and the algorithm does not require knowing the exact sparsity. Given query access to the generated evolution $\mathcal{E}_t = e^{t\mathcal{L}}$, the goal is to output estimates of all Lindbladian coefficients $\{h_i, a_{ij}\}$ to an additive error of ε with high success probability.

CONTRIBUTIONS

Our results can be broken down into two parts, summarized in [figure 1](#).

(I) Heisenberg-limited learning. Our main result is that, despite not knowing either the Hamiltonian or dissipator supports in advance, one can learn the Hamiltonian component outside the dissipator footprint with Heisenberg-limited total evolution time. To make this statement precise, we introduce the following structural objects. For a threshold η , let $\mathcal{S}_{D,\eta} := \{P_k \in \mathcal{S}_D : a_{kk} \geq \eta\}$ denote the η -heavy dissipator structure, and define its dissipator footprint by

$$\mathfrak{D}_\eta := \mathcal{S}_{D,\eta} \cup \{P = P_\ell P_m : P_\ell, P_m \in \mathcal{S}_{D,\eta}, \ell \neq m\}, \quad (2)$$

where pairwise Pauli products are taken modulo phase. As shown in [figure 1\(d\)](#), we define the Hamiltonian-Disjoint-from-Dissipator (HDD) Pauli set as $\mathcal{S}_{\text{HDD},\varepsilon} := \mathcal{S}_H \setminus \mathfrak{D}_{\eta(\varepsilon)}$, where $\eta = \eta(\varepsilon)$ is chosen as a function of the target accuracy ε . Intuitively, as the target accuracy ε decreases, weaker dissipator components can bias the Hamiltonian of interest. The algorithm must therefore resolve the dissipator down to a smaller threshold $\eta(\varepsilon)$, enlarging the dissipator footprint and shrinking the Hamiltonian target accessible at Heisenberg scaling.

Result 1 (Ansatz-free HDD Hamiltonian Learning). *There exists a quantum algorithm that, under the balanced Kossakowski tail condition (Eq. (4)), outputs an estimator of the true Hamiltonian-Disjoint-from-Dissipator coefficients such that with high probability:*

$$|\widehat{h}_k - h_k| \leq \varepsilon, \quad \forall P_k \in \mathcal{S}_{\text{HDD},\varepsilon}$$

The total evolution time under $e^{\mathcal{L}t}$ is Heisenberg-limited:

$$t_{\text{tot}} = \tilde{\mathcal{O}}(M^2/\varepsilon). \quad (3)$$

The algorithm employs only Clifford operations, uses $n + \mathcal{O}(\log M)$ noiseless ancillary qubits, and has classical processing cost $\tilde{\mathcal{O}}(\text{poly}(n, M, 1/\varepsilon))$.

This result assumes a balanced Kossakowski tail condition:

$$|a_{kj}| \leq |\mathcal{S}_D| \min\{a_{kk}, a_{jj}\} \quad \text{for all } P_k, P_j \in \mathcal{S}_D. \quad (4)$$

Intuitively, it rules out large off-diagonal Kossakowski coherences $|a_{kj}|$ between Pauli errors with highly imbalanced diagonal rates a_{kk}, a_{jj} . This condition is satisfied by many common noise models, including local and global depolarizing noise and amplitude damping as well as collective drive fluctuations of the type considered in [11].

To attain the Heisenberg limit in HDD Hamiltonian learning, we design a Lindbladian reshaping procedure based on stabilizer QEC. This procedure can be used to suppress selected Lindbladian terms, thereby enabling the long coherent evolutions needed for attaining the HL. If the structures of both the Hamiltonian and the dissipator were known a priori, one could apply the reshaping procedure directly: choose a code that suppresses the dissipative terms while allowing the relevant Hamiltonian terms to act as logical operators [6, 12, 13]. This would reduce the task to Hamiltonian learning within a logical subspace, where one could apply their favorite Heisenberg-limited Hamiltonian learning method [14–17].

To remove the known-structure assumption, we introduce a QEC-based recursive framework that underlies both the dissipator-structure learner and the HDD Hamiltonian learner highlighted in figure 1 (a). At each stage, the algorithm identifies the strongest remaining Lindbladian terms, suppresses them by adding them to the target error-correction or error-detection set, and then proceeds to weaker terms, as illustrated in figure 1 (b)–(c). By randomizing the stabilizer-code construction, we ensure that the unknown weaker terms survive as non-trivial logical operators with non-negligible probability, while the targeted terms are deterministically suppressed. We show that this can be achieved with only $\mathcal{O}(\log M)$ ancillas.

Crucially, the dissipator stage avoids SQL scaling by identifying only the η -heavy support needed for the subsequent QEC-code construction, rather than estimating dissipator coefficients to accuracy η . We prove that this heavy-support task can be performed in total evolution time $\tilde{\mathcal{O}}(1/\eta)$. Under the balanced Kossakowski-tail condition, it suffices to set $\eta = \mathcal{O}(\varepsilon/M)$ to learn the Hamiltonian-Disjoint-from-Dissipator to accuracy ε . Thus, combining the η -heavy dissipator learner with the logical Hamiltonian learner yields an end-to-end, Heisenberg-limited, ansatz-free algorithm for HDD Hamiltonian learning.

(II) Standard-quantum-limited learning. Our second result is an ansatz-free standard-limited algorithm for learning all coefficients $\{h_k\} \cup \{a_{kj}\}$ of a Lindbladian as in Eq. (1). The algorithm remains efficient even when the Lindbladian terms have unbounded locality. Its structure-identification step follows the HDD Hamiltonian learner above, but uses the more stringent threshold $\eta = \mathcal{O}(\varepsilon^2/M)$, which removes the balanced Kossakowski-tail assumption. After this step, we know the η -heavy dissipator structure and have learned the HDD Hamiltonian coefficients to accuracy ε (see figure 1 (a)). The only remaining unknowns are the dissipator coefficients and the Hamiltonian coefficients supported on the learned dissipator structure.

To learn these remaining coefficients, we use the Choi state of the short-time channel, obtained by applying $e^{\mathcal{L}t} \otimes I_n$ to the maximally entangled $2n$ -qubit state. In the Choi representation, the first-order expansion separates into orthogonal directions corresponding to individual Hamiltonian and dissipative coefficients. We exploit this property by explicitly constructing $2n$ -qubit low-rank observables whose short-time derivatives provide direct access to the desired coefficients. Thus, rather than assembling a large linear system, each remaining coefficient is isolated by the derivative of a single Choi expectation value. These derivatives are estimated using shadow process tomography [18–20] in combination with Chebyshev polynomial interpolation. This direct-observable construction is the main novel ingredient: compared with the black-box learner of [18], it improves the evolution-time scaling from $\tilde{\mathcal{O}}(nM^3/\varepsilon^4)$ to $\tilde{\mathcal{O}}(M/\varepsilon^2)$, and compared with general shadow process tomography approaches [19, 20], it avoids condition-number-dependent error amplification from linear-system inversion.

Result 2 (Ansatz-free Lindbladian Learning). *There exists a quantum algorithm for learning a Lindbladian as in Eq. (1) that outputs an estimator $\hat{\lambda} = (\hat{h}, \hat{a})$ of the true Hamiltonian and dissipator coefficient vector $\lambda = (h, a)$ such that:*

$$\|\hat{\lambda} - \lambda\|_\infty \leq \varepsilon \quad (5)$$

with high probability. The total evolution time under $e^{\mathcal{L}t}$ is

$$t_{\text{tot}} = \tilde{\mathcal{O}}\left(\frac{M}{\varepsilon^2} + \frac{M^2}{\varepsilon}\right). \quad (6)$$

The algorithm employs only Clifford operations, uses $n + \mathcal{O}(\log M)$ noiseless ancillary qubits, and has classical processing cost $\text{poly}(n, M, 1/\varepsilon)$.

The full manuscript is available on arXiv [21].

-
- [1] V. Gorini, A. Kossakowski, and E. C. G. Sudarshan, Completely positive dynamical semigroups of n -level systems, *Journal of Mathematical Physics* **17**, 821 (1976).
 - [2] G. Lindblad, On the generators of quantum dynamical semigroups, *Communications in mathematical physics* **48**, 119 (1976).
 - [3] S. F. Huelga, C. Macchiavello, T. Pellizzari, A. K. Ekert, M. B. Plenio, and J. I. Cirac, Improvement of frequency standards with quantum entanglement, *Physical Review Letters* **79**, 3865 (1997).
 - [4] R. Demkowicz-Dobrzański, J. Kołodyński, and M. Guţă, The elusive heisenberg limit in quantum-enhanced metrology, *Nature communications* **3**, 1063 (2012).
 - [5] E. M. Kessler, I. Lovchinsky, A. O. Sushkov, and M. D. Lukin, Quantum error correction for metrology, *Physical review letters* **112**, 150802 (2014).
 - [6] S. Zhou, M. Zhang, J. Preskill, and L. Jiang, Achieving the heisenberg limit in quantum metrology using quantum error correction, *Nature communications* **9**, 78 (2018).
 - [7] S. B. Antu and S. Zhou, Stabilizer codes for Heisenberg-limited many-body Hamiltonian estimation, *Quantum* **9**, 1766 (2025).
 - [8] D. Stilck França, L. A. Markovich, V. Dobrovitski, A. H. Werner, and J. Borregaard, Efficient and robust estimation of many-qubit hamiltonians, *Nature Communications* **15**, 311 (2024).
 - [9] D. S. França, T. Möbus, C. Rouzé, and A. H. Werner, *Learning and certification of local time-dependent quantum dynamics and noise* (2025), [arXiv:2510.08500 \[quant-ph\]](#).
 - [10] P. Ivashkov, N. Romanov, W. Gong, A. Gu, H.-Y. Hu, and S. F. Yelin, *Ansatz-free learning of lindbladian dynamics in situ* (2026), [arXiv:2603.05492 \[quant-ph\]](#).
 - [11] T. Kraft, M. K. Joshi, W. Lam, T. Olsacher, F. Kranzl, J. Franke, L. K. Joshi, R. Blatt, A. Smerzi, D. S. França, B. Vermersch, B. Kraus, C. F. Roos, and P. Zoller, *Bounded-error quantum simulation via hamiltonian and lindbladian learning* (2025), [arXiv:2511.23392 \[quant-ph\]](#).
 - [12] R. Demkowicz-Dobrzański, J. Czajkowski, and P. Sekatski, Adaptive quantum metrology under general markovian noise, *Phys. Rev. X* **7**, 041009 (2017).
 - [13] W. Górecki, S. Zhou, L. Jiang, and R. Demkowicz-Dobrzański, Optimal probes and error-correction schemes in multi-parameter quantum metrology, *Quantum* **4**, 288 (2020).
 - [14] H.-Y. Huang, Y. Tong, D. Fang, and Y. Su, Learning many-body hamiltonians with heisenberg-limited scaling, *Physical Review Letters* **130**, 200403 (2023).
 - [15] A. Bakshi, A. Liu, A. Moitra, and E. Tang, Structure learning of hamiltonians from real-time evolution, in *2024 IEEE 65th Annual Symposium on Foundations of Computer Science (FOCS)* (IEEE, 2024) pp. 1037–1050.
 - [16] H.-Y. Hu, M. Ma, W. Gong, Q. Ye, Y. Tong, S. T. Flammia, and S. F. Yelin, Ansatz-free hamiltonian learning with heisenberg-limited scaling, *arXiv preprint arXiv:2502.11900* (2025).
 - [17] M. Ma, S. T. Flammia, J. Preskill, and Y. Tong, Learning k -body hamiltonians via compressed sensing, *arXiv preprint arXiv:2410.18928* (2024).
 - [18] M. C. Caro, Learning quantum processes and hamiltonians via the pauli transfer matrix, *ACM Transactions on Quantum Computing* **5**, 1 (2024).
 - [19] R. Levy, D. Luo, and B. K. Clark, Classical shadows for quantum process tomography on near-term quantum computers, *Phys. Rev. Res.* **6**, 013029 (2024).
 - [20] J. Kunjummen, M. C. Tran, D. Carney, and J. M. Taylor, Shadow process tomography of quantum channels, *Phys. Rev. A* **107**, 042403 (2023).
 - [21] N. Romanov, P. Ivashkov, W. Gong, I. Kannan, A. Gu, H.-Y. Hu, and S. F. Yelin, Learning arbitrary lindbladians with quantum error correction, *arXiv preprint arXiv:2606.18188* (2026).