
Hamiltonian Learning using Machine Learning Models Trained with Continuous Measurements

Amit Kiran Rege^{1*}
amit.rege@colorado.edu

Kris Tucker^{1,2*}
ktucker27@gmail.com

Conor Smith³
cssmith36@unm.edu

Claire Monteleoni^{1,4}
cmontel@colorado.edu

¹University of Colorado Boulder, ²NASK Incorporated, ³University of New Mexico, ⁴INRIA Paris

Abstract

We present a machine learning approach for Hamiltonian parameter estimation in quantum systems using weak measurement data. Our model combines an LSTM-based encoder with a physics-informed decoder, incorporating a trainable correction term for unmodeled dynamics. Evaluated in both supervised and unsupervised settings, the method demonstrates high accuracy and robustness against noise. This approach provides an effective solution for quantum system identification in our two-qubit demonstration, outperforming traditional methods in parameter estimation tasks and offering a potentially scalable solution for quantum system identification.

1 Introduction

The characterization and control of quantum systems is crucial for advancing quantum technologies. As quantum devices grow more complex, system parameter estimation becomes challenging due to exponential state space growth, necessitating new, scalable approaches for quantum system identification. Recently, Machine learning (ML) has been successfully applied to problems in quantum state tomography [14], Hamiltonian learning [24], and quantum control [22]. However, existing methods often require strong measurements, which provides complete information but irreversibly collapses the quantum state, or explicit quantum state representations, limiting scalability. Weak measurements [31] offer an alternative approach - they extract only partial information, allowing the quantum state to be minimally disturbed and continuously monitored, though at the cost of increased noise. We present a novel ML-based approach for estimating Hamiltonian parameters from weak measurement records, leveraging advances in sequence modeling and neural ordinary differential equations (ODEs) [19].

Background and Related Work: Hamiltonian parameter estimation in quantum systems faces challenges due to measurement-induced disturbances. Weak measurements [31] maintain quantum superposition but introduce noise, complicating estimation without full state access. Recent ML applications in quantum physics have addressed challenges in state tomography [14, 13, 16, 36, 15, 38], evolution [38, 6], unknown dynamics learning [11, 18, 34], and device characterization [24, 22, 32, 1, 17]. Our work introduces a novel architecture combining LSTMs [28] with ODE solvers, inspired by Neural ODEs [19, 27, 6] and physics-informed neural networks [35]. We consider supervised (known parameters) and unsupervised (unknown parameters) scenarios. The supervised setting offers fast inference through a neural network forward pass, independent of state

*Equal contribution

representation [1]. The unsupervised case combines a physical model with a model-free correction term to handle unanticipated effects. We contrast our work with recent Hamiltonian learning methods [1, 26, 3, 2, 29, 33, 20, 4, 5] by using continuous measurements modeled by a stochastic master equation. Our approach applies without requiring steady state or energy eigenstate preparation, or adaptive Hamiltonian changes. It only requires parameters to have observable impacts on continuous measurements. Our unique discrepancy modeling combines physical system knowledge with model-free machine learning, enhancing accuracy and interpretability. While not currently used, adaptive techniques could potentially improve our learning process.

Contributions: We present a novel ML approach for Hamiltonian parameter estimation using weak measurement data, introducing: (1) A denoising autoencoder with an LSTM-based encoder and physics-informed decoder with trainable correction. (2) Demonstration of efficacy in supervised and unsupervised settings, showing accuracy and noise robustness. (3) Ability to handle unmodeled dynamics, demonstrated with unanticipated Lindblad dissipative terms. We demonstrate our approach on a two-qubit system as a proof of concept, achieving high accuracy in both supervised and unsupervised settings. While scaling to larger systems will face challenges including increased parameter space dimensionality and computational requirements, our architecture’s modular nature allows for potential extensions left as future work.

2 Problem Formulation

Our work addresses the challenge of estimating Hamiltonian parameters in a two-qubit quantum system using weak continuous measurements. The system is described by the Hamiltonian $H = \sum_{i=1}^2 \frac{\Omega}{2} X_i + \epsilon Z_1 Z_2$ where X_i and Z_i are Pauli operators acting on qubit i , Ω is the Rabi drive frequency, and ϵ is the strength of the two-qubit interaction. Our primary objective is to estimate Ω and ϵ from measurement records.

Our data consists of noisy voltage trajectories from weak continuous measurements, described by a stochastic master equation, reflecting partial information about the quantum state evolution. We provide a full description of the physical system under consideration and the stochastic master equation in Appendix A.

The dataset comprises of $N = 32,000$ trajectories for each distinct parameter set (Ω, ϵ) out of a total of $K = 80$ such tuples. We generate these trajectories using an Euler-Maruyama integrator to solve the stochastic differential equations governing the measurement records. To ensure our models generalize to unseen parameter values, the training set does not contain any trajectories corresponding to parameter values in the test and validation sets.

The training process utilizes averaged trajectories from groups of size d as input to the model to reduce diffusive noise. These groups are randomly selected from the set of N trajectories for each parameter tuple $\theta_k = (\Omega_k, \epsilon_k)$. Each mini-batch comprises $M = N/d$ averaged trajectories $\{\tilde{x}_{j,k}(t)\}_{j=1}^M$, where $\tilde{x}_{j,k}(t) = \frac{1}{d} \sum_{i \in I} r_{i,k}(t)$. Here, $r_{i,k}(t)$ represents the i -th measurement record for parameter set k , and I is a randomly selected subset of trajectory indices. For unsupervised learning, we use $x_k(t) = \frac{1}{N} \sum_{i=1}^N r_{i,k}(t)$ in the loss function, representing the average over all trajectories for a given parameter set. This approach can be interpreted as a denoising process, where $x_k(t)$ approximates the uncorrupted data, and the smaller trajectory groups represent corrupted inputs.

We consider a range of group sizes d in our experiments. We use $d = \infty$ to denote the case where all N trajectories are averaged into a single group, which approaches the noiseless measurement limit as N becomes large. Henceforth, this is called the “noise-free” setting in this paper. We describe the dataset generation and setup process in detail in Appendix B.

3 Methods and Results

Our approach includes: a supervised learning model for direct parameter estimation, and an unsupervised learning model with a physics-informed decoder when ground truth values are unknown. All results are reported averaged over 100 runs with details of training discussed in Appendix B and C.




Figure 1: Estimations of the parameter pair $\theta_k = (\Omega_k, \epsilon_k)$ averaged over four noisy trajectory groups of $d = 4,000$ measurement records for a model trained on clean data and noisy data, then evaluated on noisy data. Test set MSEs of the models are $(8.67\text{e-}3, 5.83\text{e-}4)$ and $(4.94\text{e-}3, 4.87\text{e-}4)$ respectively.

3.1 Supervised Learning

We begin with a supervised learning approach, designed to estimate Hamiltonian parameters directly from weak measurement records. This is in contrast to prior work where parameters are estimated through a post-processing step [24]. By bypassing the need for explicit quantum state representations during inference, this approach reduces memory requirements compared to full state tomography methods, though scaling to larger systems will require addressing challenges in parameter space dimensionality and measurement complexity.

Our model consists of an encoder $E : \tilde{x}(t) \rightarrow \hat{\theta}$ that maps noisy measurement trajectories to estimated parameter values. It is composed of an average pooling operation to smooth voltage trends, followed by a long short-term memory (LSTM) layer, and finally a feed-forward neural network (see Fig. 2(ii)).

The supervised model is trained by minimizing the mean squared error (MSE) between predicted and true parameters $L_{\text{supervised}} = \frac{1}{pK} \sum_k |\hat{\theta}_k - \theta_k|_2^2$ where p is the dimensionality of θ_k , and K is the number of training examples. We show that parameters can be estimated with low error (Fig. 1).

We first study the impact of noise in the training set when evaluated on noisy test data. We consider models trained with group sizes $d = 4,000$ and $d = \infty$ (noise-free) evaluated on a test set of noisy, averaged measurement records with the same group size d as the noisy training set. We can see in Fig. 1 that for many specific pairs and for the overall MSE, the model trained on the noise-free measurements is producing less accurate estimates. This is consistent in both the Ω and ϵ errors. This suggests that the model learns to account for noise, as expected, and training with noise on the level of measurements used for prediction is beneficial. We provide a plot of MSE values in Appendix D.1.

Next, Table 1 shows effect of group sizes d on the MSE for (Ω, ϵ) estimates. We observe that accuracy improves with larger group sizes, with the best performance achieved using noise-free data ($d = \infty$). We provide a comparison with non-neural network approaches in Appendix E.

3.2 Unsupervised Learning

Labeled training data may not always be available in experimental settings so we introduce an autoencoder-based unsupervised learning approach that combines a physics-informed decoder with a model-free correction term. This hybrid approach can operate without known parameter values while maintaining physical interpretability.

Our unsupervised model extends the previous architecture by adding a decoder $D : \hat{\theta} \rightarrow \hat{x}(t)$. This decoder reconstructs the measurement record by solving the unconditioned master equation using an enhanced numerical ODE integrator. The integrator is implemented as a RNN with a custom cell that

d	Best MSE	Median MSE	Mean MSE
2,000	9.87e-3, 8.76e-4	9.72e-3, 8.71e-4	9.75e-3, 9.79e-4
4,000	4.94e-3, 4.87e-4	5.23e-3, 5.44e-4	5.62e-3, 5.75e-4
8,000	4.17e-3, 2.74e-4	2.94e-3, 2.74e-4	3.38e-3, 4.13e-4
16,000	2.19e-3, 2.06e-4	1.40e-3, 1.92e-4	1.80e-3, 2.94e-4
∞	1.01e-5, 2.04e-5	4.51e-6, 4.82e-5	1.32e-5, 5.68e-5

Table 1: Each cell shows the MSE pair for (Ω, ϵ) estimates on supervised training set using the encoder only.

d	$\Delta t = 2^{-8} \mu\text{s}$	2^{-6}	2^{-4}
2,000	1.12e-2, 8.01e-4	1.10e-2, 8.42e-4	1.23e-2, 2.12e-3
4,000	5.87e-3, 4.56e-4	6.46e-3, 5.07e-4	6.39e-3, 1.68e-3
8,000	3.08e-3, 2.32e-4	2.98e-3, 2.87e-4	3.34e-3, 1.46e-3
16,000	1.68e-3, 1.19e-4	1.11e-3, 1.63e-4	1.58e-3, 1.39e-3
∞	1.25e-5, 7.75e-6	1.76e-5, 7.07e-5	1.14e-4, 1.23e-3

Table 2: MSE of (Ω, ϵ) estimates on unsupervised training set with varying Ω and ϵ and $N = 32,000$. The values of κ and η are assumed to be known exactly.

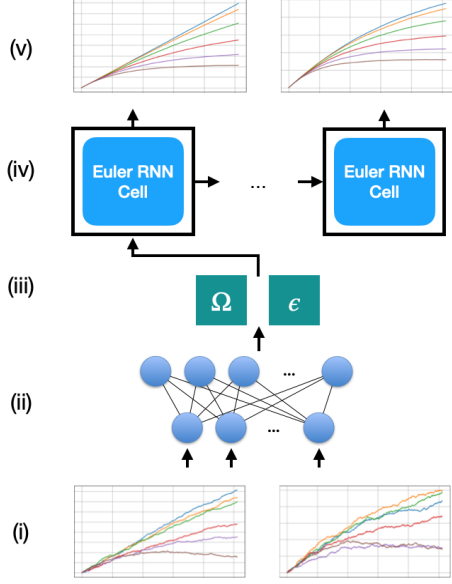


Figure 2: Diagram of the full ML model. Voltage records are measured from the cavity system producing noisy averaged voltage records (i), these are sent through a neural network encoder (ii) producing system parameters as output (iii). These parameters are used by the integrator decoder (iv) to produce noise-free voltage estimates (v).




Figure 3: Details of the decoder RNN cell combining a standard ODE integrator with an LSTM update.

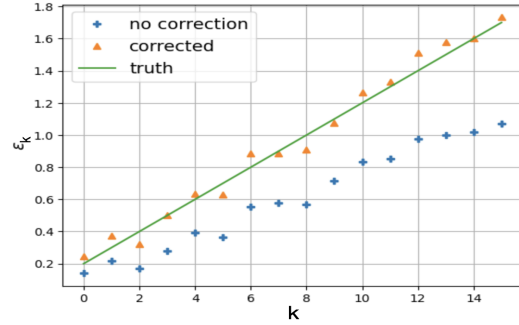


Figure 4: Estimate of ϵ for training data simulated with single particle relaxation rate $\gamma_s = 0.1$ but a model that does not explicitly account for γ_s . Here, k indexes different parameter sets tested. Results are shown for the physical model alone and with a correction learned by the decoder LSTM.

combines a single step of Euler’s method with a correction produced by a standard LSTM cell. This is shown in Fig. 3. The correction term is designed to compensate for unanticipated dynamics in the physical model used by the numerical integrator. The design leverages the concept of ODE-RNN [27, 37], as it uses a neural network to approximate corrections to the unconditioned master equation as $\frac{d\rho}{dt} = f_{\text{model}}(t, \rho; \hat{\theta}) + f_{\text{LSTM}}(t, \rho; \hat{\theta})$. Here, ρ is the density matrix, f_{model} represents the drift term derived from the master equation, and f_{LSTM} is a trainable correction term produced by an LSTM cell which is a non-linear function of the state, time and parameters. Fig. 2 shows an end-to-end illustration of our model. This architecture can be viewed as a denoising autoencoder, where the latent space corresponds to the physical parameters being estimated.

The loss function for unsupervised learning is the MSE between the decoder’s output $\hat{x}(t)$ and the provided approximations of clean measurements $x(t)$ is $L_{\text{unsupervised}} = \frac{1}{MKN_t} \sum_{j,k,t} |x_k(t) - D(E(\tilde{x}_{j,k}))(t)|^2$ where M is the number of trajectory groups per parameter set, K is the number of parameter sets, and N_t is the number of time points.

We evaluate the unsupervised model in three scenarios: varying both Ω and ϵ , fixed Ω with varying ϵ , and vice-versa. For now, we assume that the physical model in the decoder is correct, i.e. we are not learning parameters for the drift function correction f_{LSTM} . Table 2 shows results when both parameters vary. This is a fairly difficult setting as training data for each unique pair (Ω, ϵ) is small. We present ablation results on dataset size, group size, parameter and time-step in Appendix D.2.

We find that accuracy in Table 2 improves linearly with larger group sizes (doubling group size, halves MSE) and smaller time steps, with a threshold around $\Delta t = 2^{-6} \mu\text{s}$ above which this trends seems to break down suggesting the need for frequent measurements in the unsupervised case.

Model Correction for Unmodeled Dynamics: A key advantage of our unsupervised approach is its ability to correct for unmodeled dynamics. To demonstrate this, we introduce a single-particle relaxation term in the data generation ((Eq. (1a))), but not in the decoder’s physical model. We

account for this by enabling training in the decoder LSTM parameters. We provide details of our experimental setup in Appendix D.3.

Fig. 4 shows the estimation of ϵ for a ground truth value $\epsilon = 1.7$, comparing results with and without the learned correction. We observe that the learned correction significantly reduces the bias in ϵ estimation introduced by the unmodeled relaxation, demonstrating the model’s robustness to unanticipated dynamics. This capability is crucial for real-world applications where the full dynamics of the system may not be known a priori. We provide additional results estimating probabilities and MSEs in Appendix D.3.

4 Conclusion

We have proposed a machine learning model based on a denoising autoencoder capable of direct estimation of physical parameters in a two-qubit system modeled by a stochastic master equation from weak measurement records. The model is capable of learning in a supervised and an unsupervised context, and it can accurately predict parameters for systems not seen in the training data. By leveraging a master equation integrator for unsupervised learning, the autoencoder demonstrates robustness to unanticipated dynamics not included in its physical model, as shown in the case of unexpected Lindblad dissipative terms. Looking ahead, promising avenues for future research include investigating the model’s resilience to non-Markovian and nonlinear dynamics, and exploring extensions to larger quantum systems with more complex Hamiltonians - though this would require addressing challenges in parameter space dimensionality and measurement complexity.

5 Acknowledgments

We would like to thank Tameem Albash for his work and guidance on the journal version of this paper. This material is based upon work supported by the National Science Foundation the Quantum Leap Big Idea under Grant No. OMA-1936388. AKR and CM acknowledge support from NSF QLCI OMA-2016244.

References

- [1] Zheng An et al. “Unified quantum state tomography and Hamiltonian learning: A language-translation-like approach for quantum systems”. In: *Phys. Rev. Appl.* 21 (1 Jan. 2024), p. 014037. DOI: 10.1103/PhysRevApplied.21.014037. URL: <https://link.aps.org/doi/10.1103/PhysRevApplied.21.014037>.
- [2] Anurag Anshu and Srinivasan Arunachalam. “A survey on the complexity of learning quantum states”. In: *Nature Reviews Physics* 6.1 (Jan. 2024), pp. 59–69. ISSN: 2522-5820. DOI: 10.1038/s42254-023-00662-4. URL: <https://doi.org/10.1038/s42254-023-00662-4>.
- [3] Anurag Anshu et al. “Sample-efficient learning of interacting quantum systems”. In: *Nature Physics* 17.8 (Aug. 2021), pp. 931–935. ISSN: 1745-2481. DOI: 10.1038/s41567-021-01232-0. URL: <https://doi.org/10.1038/s41567-021-01232-0>.
- [4] Eyal Bairey, Itai Arad, and Netanel H. Lindner. “Learning a Local Hamiltonian from Local Measurements”. In: *Phys. Rev. Lett.* 122 (2 Jan. 2019), p. 020504. DOI: 10.1103/PhysRevLett.122.020504. URL: <https://link.aps.org/doi/10.1103/PhysRevLett.122.020504>.
- [5] Eyal Bairey et al. “Learning the dynamics of open quantum systems from their steady states”. In: *New Journal of Physics* 22.3 (Mar. 2020), p. 032001. DOI: 10.1088/1367-2630/ab73cd. URL: <https://dx.doi.org/10.1088/1367-2630/ab73cd>.
- [6] Leonardo Banchi et al. “Modelling non-markovian quantum processes with recurrent neural networks”. In: *New Journal of Physics* 20.12 (Dec. 2018), p. 123030. DOI: 10.1088/1367-2630/aaf749. URL: <https://dx.doi.org/10.1088/1367-2630/aaf749>.
- [7] Alexandre Blais et al. “Circuit quantum electrodynamics”. In: *Rev. Mod. Phys.* 93 (2 May 2021), p. 025005. DOI: 10.1103/RevModPhys.93.025005. URL: <https://link.aps.org/doi/10.1103/RevModPhys.93.025005>.

- [8] R. Bonifacio, P. Schwendimann, and Fritz Haake. “Quantum Statistical Theory of Superradiance. I”. In: *Phys. Rev. A* 4 (1 July 1971), pp. 302–313. DOI: 10.1103/PhysRevA.4.302. URL: <https://link.aps.org/doi/10.1103/PhysRevA.4.302>.
- [9] Leo Breiman. “Random Forests”. In: *Machine Learning* 45.1 (2001), pp. 5–32. ISSN: 0885-6125. DOI: 10.1023/A:1010933404324. URL: <http://dx.doi.org/10.1023/A%3A1010933404324>.
- [10] Todd A. Brun. “A simple model of quantum trajectories”. In: *American Journal of Physics* 70.7 (July 2002), pp. 719–737. ISSN: 0002-9505. DOI: 10.1119/1.1475328. URL: <https://doi.org/10.1119/1.1475328>.
- [11] Steven L. Brunton, Joshua L. Proctor, and J. Nathan Kutz. “Discovering governing equations from data by sparse identification of nonlinear dynamical systems”. In: *Proceedings of the National Academy of Sciences* 113.15 (2016), pp. 3932–3937. DOI: 10.1073/pnas.1517384113. URL: <https://www.pnas.org/doi/abs/10.1073/pnas.1517384113>.
- [12] C. C. Bultink et al. “General method for extracting the quantum efficiency of dispersive qubit readout in circuit QED”. In: *Applied Physics Letters* 112.9 (Mar. 2018), p. 092601. ISSN: 0003-6951. DOI: 10.1063/1.5015954. eprint: https://pubs.aip.org/aip/apl/article-pdf/doi/10.1063/1.5015954/13033954/092601%5B1%5D_online.pdf. URL: <https://doi.org/10.1063/1.5015954>.
- [13] Giuseppe Carleo, Yusuke Nomura, and Masatoshi Imada. “Constructing exact representations of quantum many-body systems with deep neural networks”. In: *Nature Communications* 9.1 (2018), p. 5322. DOI: 10.1038/s41467-018-07520-3. URL: <https://doi.org/10.1038/s41467-018-07520-3>.
- [14] Giuseppe Carleo and Matthias Troyer. “Solving the quantum many-body problem with artificial neural networks”. In: *Science* 355.6325 (2017), pp. 602–606. DOI: 10.1126/science.aag2302. URL: <https://www.science.org/doi/abs/10.1126/science.aag2302>.
- [15] Juan Carrasquilla and Giacomo Torlai. “How To Use Neural Networks To Investigate Quantum Many-Body Physics”. In: *PRX Quantum* 2 (4 Nov. 2021), p. 040201. DOI: 10.1103/PRXQuantum.2.040201. URL: <https://link.aps.org/doi/10.1103/PRXQuantum.2.040201>.
- [16] Juan Carrasquilla et al. “Reconstructing quantum states with generative models”. In: *Nature Machine Intelligence* 1.3 (2019), pp. 155–161. DOI: 10.1038/s42256-019-0028-1. URL: <https://doi.org/10.1038/s42256-019-0028-1>.
- [17] Leonardo K. Castelano et al. *Physics informed neural networks learning a two-qubit Hamiltonian*. 2023. arXiv: 2310.15148 [quant-ph].
- [18] Kathleen Champion et al. “Data-driven discovery of coordinates and governing equations”. In: *Proceedings of the National Academy of Sciences* 116.45 (2019), pp. 22445–22451. DOI: 10.1073/pnas.1906995116. URL: <https://www.pnas.org/doi/abs/10.1073/pnas.1906995116>.
- [19] Ricky T. Q. Chen et al. “Neural Ordinary Differential Equations”. In: *Advances in Neural Information Processing Systems*. Ed. by S. Bengio et al. Vol. 31. Curran Associates, Inc., 2018. URL: https://proceedings.neurips.cc/paper_files/paper/2018/file/69386f6bb1dfed68692a24c8686939b9-Paper.pdf.
- [20] Alicja Dutkiewicz, Thomas E. O’Brien, and Thomas Schuster. *The advantage of quantum control in many-body Hamiltonian learning*. 2023. arXiv: 2304.07172 [quant-ph]. URL: <https://arxiv.org/abs/2304.07172>.
- [21] Manuel Fernández-Delgado et al. “Do we Need Hundreds of Classifiers to Solve Real World Classification Problems?” In: *Journal of Machine Learning Research* 15.90 (2014), pp. 3133–3181. URL: <http://jmlr.org/papers/v15/delgado14a.html>.
- [22] E. Flurin et al. “Using a Recurrent Neural Network to Reconstruct Quantum Dynamics of a Superconducting Qubit from Physical Observations”. In: *Phys. Rev. X* 10 (1 Jan. 2020), p. 011006. DOI: 10.1103/PhysRevX.10.011006. URL: <https://link.aps.org/doi/10.1103/PhysRevX.10.011006>.
- [23] Jay Gambetta et al. “Quantum trajectory approach to circuit QED: Quantum jumps and the Zeno effect”. In: *Phys. Rev. A* 77 (1 Jan. 2008), p. 012112. DOI: 10.1103/PhysRevA.77.012112. URL: <https://link.aps.org/doi/10.1103/PhysRevA.77.012112>.

- [24] Élie Genois et al. “Quantum-Tailored Machine-Learning Characterization of a Superconducting Qubit”. In: *PRX Quantum* 2 (4 Dec. 2021), p. 040355. DOI: 10.1103/PRXQuantum.2.040355. URL: <https://link.aps.org/doi/10.1103/PRXQuantum.2.040355>.
- [25] Ian Goodfellow, Yoshua Bengio, and Aaron Courville. *Deep Learning*. <http://www.deeplearningbook.org>. MIT Press, 2016.
- [26] Jeongwan Haah, Robin Kothari, and Ewin Tang. “Learning quantum Hamiltonians from high-temperature Gibbs states and real-time evolutions”. In: *Nature Physics* 20.6 (June 2024), pp. 1027–1031. ISSN: 1745-2481. DOI: 10.1038/s41567-023-02376-x. URL: <https://doi.org/10.1038/s41567-023-02376-x>.
- [27] Mansura Habiba and Barak A. Pearlmutter. “Neural Ordinary Differential Equation based Recurrent Neural Network Model”. In: *2020 31st Irish Signals and Systems Conference (ISSC)*. 2020, pp. 1–6. DOI: 10.1109/ISSC49989.2020.9180182.
- [28] Sepp Hochreiter and Jürgen Schmidhuber. “Long Short-Term Memory”. In: *Neural Computation* 9.8 (1997), pp. 1735–1780.
- [29] Hsin-Yuan Huang et al. “Learning Many-Body Hamiltonians with Heisenberg-Limited Scaling”. In: *Phys. Rev. Lett.* 130 (20 May 2023), p. 200403. DOI: 10.1103/PhysRevLett.130.200403. URL: <https://link.aps.org/doi/10.1103/PhysRevLett.130.200403>.
- [30] Kurt Jacobs and Daniel A. Steck. “A straightforward introduction to continuous quantum measurement”. In: *Contemporary Physics* 47.5 (2006), pp. 279–303. DOI: 10.1080/00107510601101934. URL: <https://doi.org/10.1080/00107510601101934>.
- [31] R. E. Kastner. “Demystifying Weak Measurements”. In: *Foundations of Physics* 47.5 (Apr. 2017), pp. 697–707. ISSN: 1572-9516. DOI: 10.1007/s10701-017-0085-4. URL: <http://dx.doi.org/10.1007/s10701-017-0085-4>.
- [32] G. Koolstra et al. “Monitoring Fast Superconducting Qubit Dynamics Using a Neural Network”. In: *Phys. Rev. X* 12 (3 July 2022), p. 031017. DOI: 10.1103/PhysRevX.12.031017. URL: <https://link.aps.org/doi/10.1103/PhysRevX.12.031017>.
- [33] Haoya Li et al. *Heisenberg-limited Hamiltonian learning for interacting bosons*. 2023. arXiv: 2307.04690 [quant-ph]. URL: <https://arxiv.org/abs/2307.04690>.
- [34] Bethany Lusch, J. Nathan Kutz, and Steven L. Brunton. “Deep learning for universal linear embeddings of nonlinear dynamics”. In: *Nature Communications* 9.1 (2018), p. 4950. DOI: 10.1038/s41467-018-07210-0. URL: <https://doi.org/10.1038/s41467-018-07210-0>.
- [35] Maziar Raissi, Paris Perdikaris, and George Em Karniadakis. *Physics Informed Deep Learning (Part II): Data-driven Discovery of Nonlinear Partial Differential Equations*. 2017. arXiv: 1711.10566 [cs.AI].
- [36] Andrea Rocchetto et al. “Learning hard quantum distributions with variational autoencoders”. In: *npj Quantum Information* 4.1 (2018), p. 28. DOI: 10.1038/s41534-018-0077-z. URL: <https://doi.org/10.1038/s41534-018-0077-z>.
- [37] Yulia Rubanova, Ricky T. Q. Chen, and David K Duvenaud. “Latent Ordinary Differential Equations for Irregularly-Sampled Time Series”. In: *Advances in Neural Information Processing Systems*. Ed. by H. Wallach et al. Vol. 32. Curran Associates, Inc., 2019. URL: https://proceedings.neurips.cc/paper_files/paper/2019/file/42a6845a557bef704ad8ac9cb4461d43-Paper.pdf.
- [38] Giacomo Torlai et al. “Quantum process tomography with unsupervised learning and tensor networks”. In: *Nature Communications* 14.1 (2023), p. 2858. DOI: 10.1038/s41467-023-38332-9. URL: <https://doi.org/10.1038/s41467-023-38332-9>.

A Physical System

Our physical system of interest is that of two qubits with fixed position in a microwave cavity as illustrated in Fig. 5. A single common mode of the cavity is coupled to the computational degree of freedom of the qubits, and they are coherently driven on resonance with a Rabi drive of frequency Ω . A two-qubit interaction term is present with magnitude ϵ . A weak measurement tone [30, 10] is applied to the cavity to probe the qubit state in one of the $\{X, Y, Z\}$ directions for each qubit, with a measurement back-action dephasing rate of κ . Upon adiabatic elimination of the cavity mode [8], which we assume has dynamics evolving at a rate much faster than time scales relevant to the qubits, we can describe the system with the stochastic master equation (SME) [7, 23].

$$d\rho = -i[H, \rho]dt + \sum_{i=1}^2 \mathcal{D}[L_i](\rho)dt + \sum_{i=1}^2 \sqrt{\frac{\eta}{2}} \mathcal{H}[L_i](\rho)dW_t^{(i)}, \quad (1a)$$

$$H = \sum_{i=1}^2 \frac{\Omega}{2} X_i + \epsilon Z_1 Z_2, \quad (1b)$$

where $\{X_i, Y_i, Z_i\}$ is the set of Pauli operators for qubit $i \in \{1, 2\}$, $\mathcal{D}[L](\rho) = L\rho L^\dagger - \frac{1}{2}\{L^\dagger L\rho + \rho L^\dagger L\}$ is the Lindblad super-operator, $\mathcal{H}[L](\rho) = L\rho + \rho L^\dagger - \rho \text{Tr}[\rho(L + L^\dagger)]$ is the measurement super-operator, η is the efficiency of the measurement, and $L_i = \sqrt{\kappa}C_i$, where $C_i \in \{X_i, Y_i, Z_i\}$ is the weak measurement operator. Here we have suppressed the dependence on t for ρ , and we have adopted the normalization convention $\hbar = 1$.

The stochastic differential equations for the measurement records are given by

$$dr_i = \sqrt{\frac{\eta}{2}} \text{Tr}[\rho(L_i + L_i^\dagger)] dt + dW_t^{(i)} \quad (2)$$

where r_i is the weak measurement record for qubit $i \in \{1, 2\}$ and the independent Wiener increments $dW_t^{(i)}$ are the same as those appearing in Eq. (1a). This is comparable to the system studied in Ref. [24], except generalized to two qubits and with the addition of the two-qubit interaction term in Eq. (1b). We also consider cases where the weak measurement operator L_i is different for each qubit.

The unconditioned master equation is given by averaging over all possible trajectories of the Wiener processes and is given by:

$$d\bar{\rho} = -i[H, \bar{\rho}]dt + \sum_{i=1}^2 \mathcal{D}[L_i](\bar{\rho})dt \quad (3a)$$

$$d\bar{r}_i = \sqrt{\frac{\eta}{2}} \text{Tr}[\bar{\rho}(L_i + L_i^\dagger)] dt \quad (3b)$$




Figure 5: The physical system of two qubits in a microwave cavity subject to a Rabi drive Ω , single particle relaxation with rate γ_s , and dephasing due to weak measurement back-action with rate κ .

B Dataset and Training Details

In order to generate the data for training, validation, and testing, we use an Euler-Maruyama integrator to solve Eq. (2) for both qubits for each of 40 values of ϵ evenly spaced on $[0, 2)$ with fixed $\Omega = 1.395$

radians/ μs , and the same number of Ω values evenly spaced on $[1, 5)$ with fixed $\epsilon = 1.0$ radians/ μs for a total of $K = 80$ (Ω, ϵ) pairs. For each pair of values (Ω_k, ϵ_k), $N = 32,000$ measurement trajectories are simulated for $T = 4\mu\text{s}$. Half of the collection of pair of values (Ω_k, ϵ_k), corresponding to every other parameter pair, is used for training. The other half, corresponding to parameters midway between training values, is split evenly to be used for validation and testing. This ensures the training set contains a disjoint set of parameter pairs from the validation/test set. The two endpoints of the validation/test set are excluded from each end to ensure that the training data extends slightly beyond the domain of validation/test values. Only weak measurement records are being used, and no strong measurements are being performed.

Values of $\kappa = 3.326$ radians/ μs and $\eta = 0.1469$ are selected to be consistent with the example of superconducting qubits as found, for example, in Ref. [24], though with a stronger measurement back-action κ after simulations showed a larger value to offer a good balance between dephasing rate and measurement noise. While either Ω, ϵ , or both are treated as unknown parameters and comprise the latent space of our model, it is assumed in our simulations that κ and η are known with high precision as they have already been calibrated. Our approach can be used to perform this calibration step, since the latent space of the model is not limited to Hamiltonian parameters and could support any combination of Hamiltonian and Lindblad parameters as unknown variables. Alternatively, other calibration methods found in the literature can be used to estimate them, such as in [12]. It should be noted, however, that in practice if calibration errors are present in the model values, the degrees of freedom in the decoder should make parameter estimation robust to these errors as demonstrated in what is likely the more challenging case of unanticipated single particle relaxation examined in Sec. 3.2. The impact on parameter estimation accuracy is expected to be small for small calibration errors. A study quantifying this dependency is an interesting area for future work.

During training, for each true parameter set $\theta_k = (\Omega_k, \epsilon_k)$, trajectory groups of a preset size d are randomly selected from the full training set and their averages are provided as input values to the model, such that each mini-batch is comprised of $M = N/d$ averaged trajectories $\{\tilde{x}_{j,k}(t)\}_{j=1}^M$ where $\tilde{x}_{j,k}(t) = \frac{1}{d} \sum_{i \in I} r_{i,k}(t)$, $r_{i,k}(t)$ is measurement record i for parameter set k , and I is a set of trajectory indices of size d randomly selected from $1, 2, \dots, N$ without replacement until all N trajectories are used, which defines one epoch of training. In this way, each mini-batch consists of a different set of noisy trajectories as input to maximize the diversity of training examples. In the unsupervised case, for all $j = 1, \dots, M$, the value $x_k(t)$ to be used in the loss function is the average over the full training set of trajectories associated with the true parameter set θ_k , such that $x_k(t) = \frac{1}{N} \sum_{i=1}^N r_{i,k}(t)$.

To make the interpretation of our model as a denoising autoencoder more concrete, we note that DNAs are characterized by a corruption process $C(\tilde{y}|y)$ whereby noisy inputs \tilde{y} are generated for each uncorrupted training example y [25]. In our case, the full trajectory means $x_k(t)$ take on the role of y , approximating the solution to the unconditioned SME, while the random selection of much smaller trajectory groups is the process by which the corrupted data elements \tilde{y} are generated.

C Evaluation Details

During evaluation, validation and test error is evaluated in a similar manner to how training mini-batches are selected, with groups of size d randomly chosen from the full validation and test sets, followed by the calculation of the MSE. This process is repeated 100 times and the average is taken as the calculated MSE for each set.

Training is performed for multiple runs of 100 epochs, with a learning rate of 3×10^{-3} and a decay rate of 0.99 per epoch, resetting the learning rate after each run, continuing until the validation loss drops by less than five percent run-to-run, and for the last 20 epochs of the final run. Once converged, parameter MSE is computed for the validation and test sets. This entire process is performed 100 times, each with a different random initialization of the model parameters. The best model is considered to be the one with the smallest validation loss after the final epoch, and this model is used to evaluate the MSE for the test set. Hyperparameter tuning for the model layer sizes was performed using a grid search for a physical system with parameters distinct from those in Sec. 3.

During evaluation, we consider a range of trajectory group sizes d as well as noise-free data derived from an Euler integration of the unconditioned master equation. The noise-free case is denoted by ∞ as the number of trajectories in all tables.

To minimize the impact of the weak measurement back-action via the parameter κ , when evaluating in the supervised and unsupervised settings, the initial state is chosen to be spin-up in the directions of measurement, and the directions of X and Y were used for the first and second qubit, respectively. A different configuration is used when we are testing the model correction, where more diverse measurements are needed to correct for unanticipated single particle relaxation not explicitly present in the decoder’s physical model.

D Additional Results

In this section, we report additional results from our experiments not included in the main text due to space constraints.

D.1 Supervised Setting

We provide a plot to visualize error values for the experiment on clean vs noisy training data performed in Sec. 3.1. It can clearly be seen that the model trained on noisy data (orange) achieves lower error compared to the one trained on noise-free data.




Figure 6: Estimations of the squared error averaged over four noisy trajectory groups of $d = 4,000$ measurement records for a model trained on clean data and noisy data, then evaluated on noisy data. Shuffle evaluated test set MSEs are $(8.67e-3, 5.83e-4)$ and $(4.94e-3, 4.87e-4)$ respectively. Here, k indexes different parameter sets tested

We further compare our supervised approach to a non-neural network baseline in Appendix E.

D.2 Unsupervised Setting

We now present ablation studies over specific choices made for the experiments and results presented in Sec. 3.2.

First we examine how parameter estimation accuracy for a fixed group size $d = 4,000$ varies with the total training set size N . Table 3 shows (Ω, ϵ) test set MSEs for the case where Ω is unknown but fixed at a true value of 1.395 radians/ μ s and ϵ is allowed to vary, for various training set sizes N . The same test set containing 16,000 trajectories is used for each row. From the table, we see that accuracy improves significantly as the amount of training data increases, even though the groups being presented to the model for parameter estimation remain the same. This indicates that a greater diversity of noisy measurement records when training results in models that can produce more accurate parameter estimates when presented with the same number of measurements when performing prediction. This motivates using $N = 32,000$ in all our experiments.

Next, we consider the impact of group size d and measurement record time spacing Δt on accuracy. Table 4 again shows (Ω, ϵ) test set MSEs for the case where Ω is unknown but fixed at a true value of 1.395 radians/ μ s and ϵ is allowed to vary. It illustrates how extremely low test set MSEs are achievable for Ω in this case, which is expected given the high volume of training data available for a single value. The accuracy of ϵ estimates depends on both the number of trajectories used to create each input sequence, as well as the time spacing at which measurements are recorded. To

N	Best MSE	Median MSE	Mean MSE
4,000	1.51e-3, 1.12e-3	6.58e-4, 9.86e-4	8.33e-4, 1.08e-3
8,000	3.27e-4, 7.71e-4	5.20e-4, 7.94e-4	5.38e-4, 8.81e-4
16,000	1.83e-4, 7.49e-4	3.20e-4, 8.41e-4	3.07e-4, 8.38e-4
32,000	8.14e-5, 7.16e-4	9.82e-5, 8.15e-4	1.45e-4, 8.12e-4

Table 3: MSE of (Ω, ϵ) when using $d = 4,000$ trajectories to estimate parameters for varying training set sizes N for an unsupervised training set with fixed Ω and varying ϵ . The values of κ and η are assumed to be known exactly.

d	$\Delta t = 2^{-8} \mu s$	2^{-6}	2^{-4}
2,000	5.78e-5, 1.40e-3	1.20e-4, 1.40e-3	1.22e-4, 2.83e-3
4,000	8.14e-5, 7.16e-4	1.66e-4, 7.90e-4	1.19e-4, 2.09e-3
8,000	8.51e-5, 3.42e-4	7.16e-5, 3.63e-4	8.02e-5, 1.82e-3
16,000	8.31e-5, 1.57e-4	1.66e-4, 1.57e-4	1.19e-4, 1.89e-3
∞	1.66e-6, 6.08e-6	7.21e-6, 1.04e-4	1.52e-5, 1.71e-3

Table 4: MSE of (Ω, ϵ) estimates on unsupervised training set with fixed Ω and varying ϵ and $N = 32,000$. The values of κ and η are assumed to be known exactly.




Figure 7: Estimate of $P(Y_i = 0)$ when $\epsilon = 1.7$ and measuring X with spin-up in Z as the initial state for both qubits for training data simulated with single particle relaxation rate $\gamma_s = 0.1$ but a model that does not explicitly account for γ_s . Results are shown for the physical model alone and with a correction learned by the decoder LSTM.

evaluate measurement records with varying time spacing, trajectories simulated with $\Delta t = 2^{-8} \mu s$ are sub-sampled to avoid introducing numerical integration error associated with simulations with a larger time step.

D.3 Model Correction

In this section, we provide experimental details and additional evidence for the ability of the decoder to correct for dynamics not explicitly considered in the physical model (Eq. (1a)).

We begin by providing details of our experimental setup. Unlike the vanilla parameter estimation setup, we now enable training for the decoder LSTM parameters as well. For the datasets, we simulate 30,000 trajectories with fixed $\Omega = 1.395$ radians/ μs that is known to the model, but with varying unknown ϵ which is the parameter to be estimated. We add a dissipative term for each qubit corresponding to the Lindblad operator $\mathcal{D}[\sqrt{\gamma_s} \sigma_i^-](\rho)$ in Eq. (1a), where $\gamma_s = 0.1$ and $\sigma_i^- = \frac{1}{2}(X_i - iY_i)$ is the single particle relaxation operator mapping the excited state to the ground state. Note that while the choice of a single unknown parameter is made in this section to provide a clear and straightforward demonstration, the model still supports multiple possible parameters to be estimated in addition to the decoder degrees of freedom.

d	$\gamma_s = 0.1$	$\gamma_s = 0.0$	$\gamma_s = 0.0 + \text{correction}$
5,000	2.84e-3	0.142	3.82e-3
∞	2.12e-6	0.145	2.71e-4

Table 5: MSE of ϵ estimates using different γ_s values in the model. The final column shows results when learning is enabled for the free parameters in the decoder to account for the unanticipated term γ_s . Measurement time spacing is $\Delta t = 2^{-8}\mu s$.

In this case, we take a more diverse set of measurements, simulating 10,000 trajectories measuring in each of the X , Y , and Z directions for both qubits, and relaxing κ to one-fourth the value used in previous setups to reduce measurement back-action. Spin-up in the Z direction is the initial state for each qubit. This more cautious approach to measurement is warranted if completely unknown effects are expected to be present. More information about the type of phenomenon, but not necessarily the magnitude, could allow for a more targeted measurement scheme, but here we keep it general. The trajectory group size used for each input was $d = 5,000$, and the best of 20 randomly initialized models was selected for the results in this section.

The results of the parameter estimation with and without the correction are shown in Fig. 7 and Table 5. Similar to the results introduced earlier in Fig. 4, we can see in Fig. 7 that single particle relaxation has introduced a significant bias to the estimated ϵ parameters when unaccounted for in the model. However, the decoder LSTM has successfully corrected for the effect, returning the MSE (Table 5) to a value much closer to where it would have been had the physical model explicitly accounted for it.

E Baseline: Random Forests

To justify our choice of employing LSTMs as the backbone of our proposed architecture, we provide a comparison using Random Forests [9] for the encoder-only experiments described in Section 3.1 with. We choose to compare against Random Forests because they have been shown to have robust off-the-shelf performance across a wide variety of tasks [21]. In this section, we give an overview of our setup and a comparison with results in Table 1, which justify the use of Neural Network based architectures in our subsequent experiments.

E.1 Training Setup

Our aim is to compare the prediction accuracy for the parameter ϵ for a fixed Ω in the case of supervised learning between the Random Forest model and the model described in the main text. To do so, we use the same procedure as described in Section 3.1 of the main text. The training set contains averaged trajectories sampled randomly in groups of varying size d , and the trajectories in the dataset correspond to 40 values of ϵ evenly spaced on $[0, 2)$ for a fixed $\Omega = 1.395$ radians/ μs . The crucial hyperparameters for our method are the number of trees in the ensemble (set to 100) and the maximum depth of each tree (set to 25). These were found via 5-fold cross validation over the dataset.

A clear point of difference compared to the LSTM-based model is the form of the input given to the Random Forest model. Unlike the LSTM model, input to the Random Forest model *must* be 1-dimensional. We use the concept of ‘dimensionality’ to refer to the number of weak measurement values at each timestep. Our input is 2-dimensional since we have a weak measurement value for each qubit. This dichotomy requires us to ‘flatten’ the input values into a single dimension. This can be done in two ways: either by keeping measurement values corresponding to a particular time step close together (column-major order) or by concatenating the values of the second qubit after *all* the values of the first qubit (row-major order). We find that the choice of ordering makes little difference in our results as we demonstrate next.

E.2 Results

We present the impact of the choice of input order in our Random Forest model in Table 6. We find the difference in MSE values to be minimal.

d	Best MSE	Median MSE	Mean MSE
2,000	2.35e-3, 2.63e-3	4.13e-3, 4.08e-3	4.41e-3, 4.52e-3
4,000	1.23e-3, 1.31e-3	3.44e-3, 3.52e-3	3.69e-3, 3.70e-3
8,000	7.49e-4, 7.34e-4	1.61e-3, 1.68e-3	1.86e-3, 1.86e-3
16,000	3.36e-4, 3.26e-4	1.47e-3, 1.50e-3	1.65e-3, 1.71e-3
∞	2.34e-4, 3.48e-5	1.54e-3, 1.37e-3	1.75e-3, 1.79e-3

Table 6: Each cell shows MSEs for Random Forest models with row-major input flattening and column-major input flattening respectively.

We now compare the results of the Random Forests model in column-major order to the LSTM model in Table 7. We find about an order of magnitude difference in the MSEs achieved by the Random Forest model compared to the LSTM model. The trend of improved accuracy with increasing sizes of d , also observed in the LSTM results, continues to hold in this case. To get a better sense of the uncertainty, we also present the mean MSEs and their one standard deviation error bars for both our models in Fig. 8. The fact that our choice of ‘flattening’ does not affect overall performance of the Random Forest model suggests that the difference in performance between the Random Forest model and the LSTM model could be attributed to the LSTM-based model taking the sequential nature of the data into account versus the Random Forest model treating each trajectory as just a high-dimensional vector. We leave a thorough investigation of these issue to future work.




Figure 8: Mean MSEs and their 1 standard deviation error bars for the Random Forests model with different input orderings and the LSTM model.

d	Best MSE	Median MSE	Mean MSE
2,000	2.63e-3, 1.47e-3	4.08e-3, 1.51e-3	4.52e-3, 1.63e-3
4,000	1.31e-3, 7.37e-4	3.52e-3, 7.16e-4	3.70e-3, 8.56e-4
8,000	7.34e-4, 4.39e-4	1.68e-3, 6.00e-4	1.86e-3, 4.67e-4
16,000	3.26e-4, 1.83e-4	1.50e-3, 2.13e-4	1.71e-3, 2.68e-4
∞	3.48e-5, 8.11e-6	1.37e-3, 5.59e-5	1.79e-3, 7.41e-5

Table 7: Each cell shows the MSE pair for ϵ estimates on supervised training set using the encoder only. The first value is the MSE of the Random Forest model using column-major ordering, and second value corresponds to the MSE of the LSTM-based model.