Monitoring Fast Superconducting Qubit Dynamics Using a Neural Network

G. Koolstra,[1,2,*] N. Stevenson,[1] S. Barzili,[3,4] L. Burns,[3,4] K. Siva,[1] S. Greenfield,[3,5] W. Livingston,[1] A. Hashim,[1,2] R. K. Naik,[1,2] J. M. Kreikebaum,[1,6] K. P. O'Brien,[7] D. I. Santiago,[1,2] J. Dressel,[3,4] and I. Siddiqi[1,2,6]

[1]*Quantum Nanoelectronics Laboratory, Department of Physics, University of California at Berkeley, Berkeley, California 94720, USA*
[2]*Computational Research Division, Lawrence Berkeley National Laboratory, Berkeley, California 94720, USA*
[3]*Institute for Quantum Studies, Chapman University, Orange, California 92866, USA*
[4]*Schmid College of Science and Technology, Chapman University, Orange, California 92866, USA*
[5]*Department of Physics and Astronomy, University of Southern California, Los Angeles, California 90089, USA*
[6]*Materials Sciences Division, Lawrence Berkeley National Laboratory, Berkeley, California 94720, USA*
[7]*Department of Electrical Engineering and Computer Science, Massachusetts Institute of Technology, Cambridge, Massachusetts 02139, USA*

(Received 6 September 2021; revised 21 April 2022; accepted 17 June 2022; published 26 July 2022)

Weak measurements of a superconducting qubit produce noisy voltage signals that are weakly correlated with the qubit state. To recover individual quantum trajectories from these noisy signals, traditional methods require slow qubit dynamics and substantial prior information in the form of calibration experiments. Monitoring rapid qubit dynamics, e.g., during quantum gates, requires more complicated methods with increased demand for prior information. Here, we experimentally demonstrate an alternative method for accurately tracking rapidly driven superconducting qubit trajectories that uses a long short-term memory (LSTM) artificial neural network with minimal prior information. Despite few training assumptions, the LSTM produces trajectories that include qubit-readout resonator correlations due to a finite detection bandwidth. In addition to revealing rotated measurement eigenstates and a reduced measurement rate in agreement with theory for a fixed drive, the trained LSTM also correctly reconstructs evolution for an unknown drive with rapid modulation. Our work enables new applications of weak measurements with faster or initially unknown qubit dynamics, such as the diagnosis of coherent errors in quantum gates.

DOI: 10.1103/PhysRevX.12.031017 Subject Areas: Quantum Physics, Quantum Information

I. INTRODUCTION

Weak measurements allow the observer to reconstruct the dynamics of a quantum system and to track the evolution of a wave function before its collapse to an eigenstate. For superconducting quantum circuits, in particular, reconstructing individual quantum trajectories [1–10] has served as a tool to monitor quantum jumps [11–14], track diffusion statistics [15–18], generate entanglement via measurement [19,20], coherently control quantum evolution using feedback [21–25], and implement continuous quantum error correction [26–29]. In weak measurements with superconducting qubits coupled to a readout resonator, the dynamics of the latter are typically much faster than the former. Therefore, the field exiting the resonator is a direct measure of the qubit dynamics, which can be readily computed using quantum filters built on generalizations of Bayes' rule [9,10], Markovian stochastic master equations [8,30], or stochastic path integrals [31]. However, such quantum filters are not well suited to cases where the qubit dynamics are necessarily fast, such as during a rapid entangling gate or syndrome measurement in an error-correction scheme [29]. Interpreting the output cavity field in this non-Markovian regime for the qubit is challenging.

A quantum filter obtained by training a long short-term memory (LSTM) neural network [32] is an attractive alternative for reconstructing trajectories in more challenging regimes. An LSTM is a variant of recurrent neural network with persistent memory, which is well suited for time series data with long temporal correlations [33]. They have recently found applications in quantum physics [34] ranging from quantum state tomography [35,36] to qubit

*gkoolstra@lbl.gov

noise spectroscopy [37,38]. The first application of an LSTM quantum filter to weak measurements successfully tracked slower qubit dynamics with a fidelity comparable to a standard quantum filter [32]. We build upon this initial success to show that LSTM-based filters are not constrained to slow qubit dynamics, but can learn rapid and nontrivial dynamics robustly without additional prior information (e.g., from calibration measurements) which would be required by traditional methods.

In this work, we demonstrate that an LSTM accurately reconstructs quantum state trajectories of a driven superconducting qubit coupled to a readout resonator, even for qubit dynamics faster than the relaxation timescale of the resonator. Our LSTM trains entirely on experimental observations, compensates for the delays and correlations originating from the limited detection bandwidth, and outperforms conventional reconstruction methods in the case of fast qubit dynamics. By extracting and separating the coherent dynamics and measurement backaction learned by the LSTM, we observe nontrivial drive-dependent corrections consistent with theory that includes the joint dynamics of the qubit and resonator (information not given to the LSTM). The dominant corrections include a rotation of the measured eigenstates and a reduction in the effective measurement rate with increasing drive. Finally, we demonstrate that the trained LSTM can correctly identify and track *a priori* unknown time dependencies in the dynamics produced by modulating the drive, despite no prior examples of such behavior in the training data. Our results suggest that LSTM-based filters may enable novel applications of continuous monitoring to previously unobtainable regimes.

Compared with the first implementation of an LSTM filter to weak measurements [32], our LSTM operates in a challenging regime, where input data are more strongly correlated, making it more difficult to interpret the physics embedded in the LSTM. Our work further extends the techniques of Ref. [32] by introducing a time-windowed analysis method for parameter estimation of time-dependent Hamiltonians. Finally, we show that the LSTM can faithfully reconstruct heterodyne measurement backaction (in contrast to homodyne measurement in Ref. [32]).

II. EXPERIMENTAL SYSTEM

Our qubit-resonator system [Fig. 1(a) and Appendixes A and B] consists of a superconducting transmon qubit (frequency $\omega_{ge}/2\pi = 5.473$ GHz) capacitively coupled to a superconducting coplanar waveguide resonator ($\omega_{\mathrm{res}}/2\pi = 6.679$ GHz). The interaction Hamiltonian in this regime is dispersive,

$$H_{\mathrm{int}} = \hbar\chi a^{\dagger}a\sigma_z, \tag{1}$$

where $\hbar$ is the reduced Planck's constant, $\chi/2\pi = 0.47$ MHz is the qubit-dependent resonator frequency shift, $a^{\dagger}$ (a) is

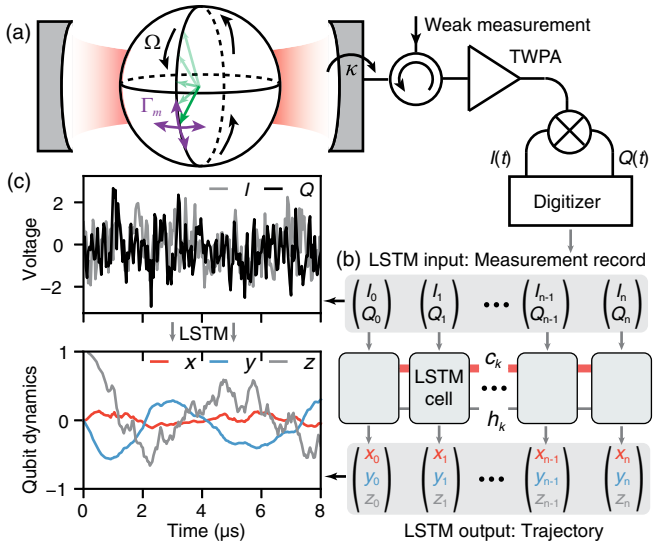

FIG. 1. Monitoring fast qubit dynamics using an LSTM artificial neural network. (a) During weak measurement of a strongly driven qubit, the qubit state (green arrow) evolves due to a coherent Rabi drive (black, rate Ω) and backaction from both quadratures of a heterodyne measurement (purple arrows). If Ω exceeds the resonator relaxation rate $\kappa/2$, memory effects such as delay and qubit-resonator correlations can no longer be ignored when reconstructing trajectories. (b) The LSTM learns the stochastic, dissipative qubit dynamics, mapping elements of each noisy voltage record (I, Q) to a qubit Bloch vector trajectory (x, y, z). The cell state c_k and hidden state h_k can encode long-term correlations in the measurement record. (c) Sample record with corresponding qubit trajectory output by the trained LSTM.

the creation (annihilation) operator for the readout resonator, and σ_z is the qubit Pauli z operator. We additionally drive the qubit (at its Stark-shifted transition frequency) to induce Rabi oscillations about the x axis at a variable Rabi frequency Ω.

To continuously monitor this system we apply a weak probe tone to the resonator at the frequency midpoint between qubit-shifted resonances, which populates the resonator with an approximate mean photon number of $\bar{n} \approx 0.3$. After interacting with the qubit, the resonator field escapes the resonator at a rate $\kappa/2 \approx 2\pi \times 0.8$ MHz. We amplify both in-phase (I) and quadrature (Q) parts of this field using a traveling wave parametric amplifier (TWPA), perform a heterodyne measurement of its complex amplitude, and digitize the resulting pair of noisy quadrature voltages with a 1 ns sampling time. After a chosen monitoring duration $0 < T_m < 8$ μs of variable length [Fig. 1(c)] we turn off the Rabi drive and projectively measure the qubit along one of the cardinal axes of the Bloch sphere ($\sigma_{x,y,z}$). We use this final projective measurement as a training label and for verification via conditioned final state tomography.

Our phase-preserving detection technique [16,39] gives rise to two distinct types of measurement backaction: non-unitary partial collapse from information that distinguishes

the qubit states, and unitary phase drift from fluctuations in resonator photon number [10]. Our detection technique is robust to drifts in amplifier gain, spans several gigahertz, and in contrast to phase-sensitive detection of just the informational quadrature [32], allows for simultaneous weak measurement of multiple qubits [19] in future experiments.

Before passing a monitoring record into a quantum filter, we digitally filter it to remove qubit-independent high-frequency noise, then coarse grain it into longer time bins while ensuring rapidly driven qubit dynamics is not undersampled. We then pass the coarse-grained digital record $\{(I_k, Q_k)\}_{k=0}^{n}$ into both the LSTM filter and a more traditional Bayesian quantum filter, and compare the resulting qubit trajectories expressed in the Bloch coordinates $\{(x_k, y_k, z_k)\}_{k=0}^{n}$.

We train the LSTM [Fig. 1(b)] to reconstruct trajectories by feeding it a subset of experimental records $\{(I_k, Q_k)\}_{k=0}^{n}$ reserved for training, which are of varying length T_m and labeled by their final projective measurement basis and outcome. During training the LSTM adjusts its network parameters to minimize the cross-entropy loss, which is computed from predictions at $t = T_m$ and associated projective measurement outcomes (for details see Appendix D). We have found that the LSTM network converges better than a simple feed forward neural network, which we attribute to the LSTM's ability to track long correlations in the cell state. After the LSTM converges we feed the trained model individual measurement records [Fig. 1(c)] from separate datasets not used for training. In the rest of this work, we analyze the trajectories $\{(x_k, y_k, z_k)\}_{k=0}^{n}$ predicted by the trained network from such test data.

III. QUBIT TRAJECTORY RECONSTRUCTION WITH A NEURAL NETWORK

Qubit dynamics faster than the resonator relaxation rate $\kappa/2$ can prevent the conditioned coherent steady states in the resonator from adiabatically following their associated qubit states [10]. To investigate the onset of this nonadiabatic behavior beyond the coherent-state approximation [8,10], we experimentally vary the rate of the qubit dynamics by tuning Rabi frequency Ω from slow dynamics ($2\Omega/\kappa \ll 1$) to fast dynamics that are well outside the adiabatic regime ($2\Omega/\kappa > 1$), training the LSTM to predict trajectories for each Rabi frequency independently.

We first use the LSTM to reconstruct qubit dynamics with a weak Rabi drive ($2\Omega/\kappa = 0.2$), where conventional steady-state methods can still accurately reconstruct trajectories. A histogram of reconstructed trajectories [Fig. 2(a)] shows oscillations due to the Rabi drive as well as diffusion due to measurement backaction, which slowly collapses the trajectories toward $|\pm Z\rangle$. The competition between (i) the Rabi drive, (ii) trajectory collapse and diffusion at rate $\eta\Gamma_m$, and (iii) trajectory dephasing at rate $2(1-\eta)\Gamma_m$ confines trajectories to a Bloch sphere with reduced radius [dashed lines in Fig. 2(a); see also Appendix E] set by the measurement efficiency η. Here, Γ_m is defined as the single quadrature measurement dephasing rate which is set by the amplitude of the weak measurement tone.

When comparing individual trajectories, the LSTM produces trajectories similar to a conventional Bayesian filter approach. This Bayesian filter method sequentially estimates the quantum state from a measurement record using a well-known initial state and calibrated values for Ω, the ensemble decay rate, and qubit-conditioned resonator output fields. We quantify the error of both methods self-consistently by averaging projective measurement results of trajectories with similar predictions [16,17,32]. The averaged tomography results closely follow the LSTM predictions for all three Bloch coordinates [Fig. 2(b)] with rms error $(\varepsilon_x^{\mathrm{LSTM}}, \varepsilon_y^{\mathrm{LSTM}}, \varepsilon_z^{\mathrm{LSTM}}) = (4.0, 4.9, 4.3) \times 10^{-2}$ [definition in Eq. (F1)]. This error largely reflects our imperfect knowledge of the true quantum state from a finite number of projective measurements; see Appendix F. More training data are likely to improve both the estimate of the true quantum state and the LSTM accuracy, since the LSTM training loss does not saturate when training on increasing fractions of available training data. We expect this saturation for a dataset size of $\mathcal{O}(10^6)$ voltage records [40]. Nevertheless, even with a finite-size training dataset considered here ($\sim 4 \times 10^5$ voltage records), the LSTM learns an accurate representation of the reduced qubit dynamics, with a lower bound on the accuracy of $1 - \varepsilon_{\mathrm{av}}^{\mathrm{LSTM}}/2 = 0.978$.

For large Rabi frequencies the LSTM trajectories remain faithful, even when the qubit dynamics exceeds the relaxation rate of the resonator $\kappa/2$ [Figs. 2(c) and 2(d)]. In contrast, the Bayesian filter's validation error increases sharply past $2\Omega/\kappa \approx 1$, because this method is sensitive to two errors related to the sampling time of voltage records, dt. The first error, due to stepwise application of fast coherent dynamics, grows with Ωdt [41], whereas the second error due to temporal correlations grows for $dt < 2/\kappa$ [42]. Consequently, for fast Rabi frequencies $2\Omega/\kappa > 1$ there is no choice of dt for which the standard Bayesian filter produces accurate trajectories.

Experimentally the breakdown of the Bayesian filter coincides with a notable delay between oscillations in the measured voltage records and the qubit coordinate $z(t)$ (see Fig. 12), which occurs because the resonator retains photons for a time κ^{-1}. In addition, the coherent oscillation amplitude in the voltage record reduces significantly for $2\Omega/\kappa > 1$, because instead of measuring the instantaneous qubit state, photons measure the time-averaged measurement operator [43]. This effect is similar to the suppression of current oscillations in a semiconducting point contact detector [44].

As a first-order correction of the resonator memory effects, we adjust the Bayesian filter based on independent master equation simulations of the resonator field

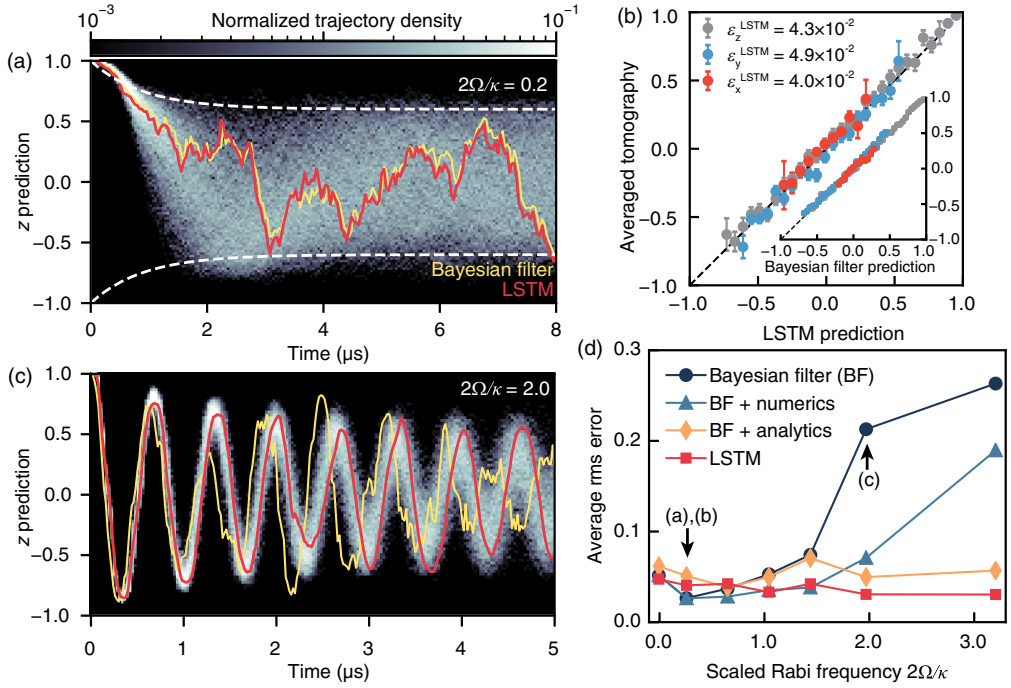

FIG. 2. Breakdown of adiabatic trajectory reconstruction. (a) The histogram of weakly driven trajectories of length $T_m = 8.0\ \mu$s (reconstructed by the LSTM with $dt = 40$ ns) shows rapid trajectory diffusion due to measurement backaction. The color bar represents the trajectory probability density at each time step. The LSTM produces trajectories (example in red) comparable to those from a steady-state Bayesian filter (yellow). (b) LSTM validation based on tomographic measurements immediately following the LSTM prediction, for the trajectories shown in (a) where $2\Omega/\kappa = 0.2$. The dashed line with slope 1 indicates perfect validation. The inset shows the Bayesian filter validation for the same dataset with slightly smaller rms error $(\varepsilon_x^{\mathrm{BF}}, \varepsilon_y^{\mathrm{BF}}, \varepsilon_z^{\mathrm{BF}}) = (2.3, 3.2, 2.5) \times 10^{-2}$. This is partly because $\varepsilon_{x,y,z}^{\mathrm{BF}}$ are computed using the entire dataset, whereas $\varepsilon_{x,y,z}^{\mathrm{LSTM}}$ only have access to trajectories not used for training. (c) For fast qubit dynamics ($2\Omega/\kappa = 2.0$) outside the adiabatic regime the predictions of the steady-state Bayesian filter and LSTM diverge (trajectory $dt = 20$ ns). (d) Validation errors averaged over the three qubit coordinates versus Rabi frequency $2\Omega/\kappa$, showing a breakdown of the steady-state Bayesian filter for $2\Omega/\kappa > 1$ (dots, BF), while the LSTM validation error (squares) remains small. The performance of the Bayesian filter improves with additional numerical prior information of the expected evolution of the z-conditioned resonator fields (BF + numerics; see Appendix G), and improves even further when adding analytical corrections to the measurement backaction (BF + analytics; see Appendix K). Importantly, the LSTM performance stays consistent without additional prior information. Arrows mark the data shown in (a)–(c).

amplitudes conditioned on the qubit states [Fig. 2(d), "BF + numerics," and further details in Appendix G]. While this reduces the rms error at moderate Rabi frequencies, the average error remains large for $2\Omega/\kappa > 1$. Only after an analytic treatment of the resonator memory, and including calibrated parameters such as χ and κ, does the Bayesian filter reach similar error as the LSTM for all Rabi frequencies [Fig. 2(d), "BF+analytics," and further details in Appendix K]. This indicates that with enough prior information the Bayesian filter with analytic correction reproduces the essential ingredients of the LSTM trajectories. However, the LSTM offers an accurate reconstruction method requiring no prior knowledge of coupling rates to the environment.

IV. RESONATOR MEMORY CORRECTIONS TO QUBIT TRAJECTORIES

To gain further insight into why the LSTM outperforms the Bayesian filter for $2\Omega/\kappa > 1$, we perform statistical analysis on the LSTM trajectories. First, we bin the LSTM trajectories inside the Bloch sphere [Fig. 3(a)] and calculate the Bloch vector increments $d\mathbf{r} = \mathbf{r}_{t+1} - \mathbf{r}_t$ for each bin, where $\mathbf{r} = (x, y, z)$. Since the Rabi drive confines most of the interesting qubit dynamics to the yz plane, we restrict our analysis to two dimensions. Therefore, each bin contains a set of Bloch vector increments $\{(dy_k, dz_k)\}_{k=0}^m$, where each pair can be seen as a random variable due to the stochastic nature of the measurement backaction [Fig. 3(b), dots]. Next, we compute the average drift $(d\bar{y}, d\bar{z})$ and diffusion from the sample mean and covariance matrix on $\{(dy_k, dz_k)\}_{k=0}^m$, respectively. The coviarance matrix,

$$C(y,z) = \begin{pmatrix} \sigma_{dy}^2 & \sigma_{dy,dz} \\ \sigma_{dy,dz} & \sigma_{dz}^2 \end{pmatrix}, \tag{2}$$

$$\sigma_{dy}^2 = \frac{1}{m}\sum_{k=0}^{m}(dy_k - d\bar{y})^2, \tag{3}$$

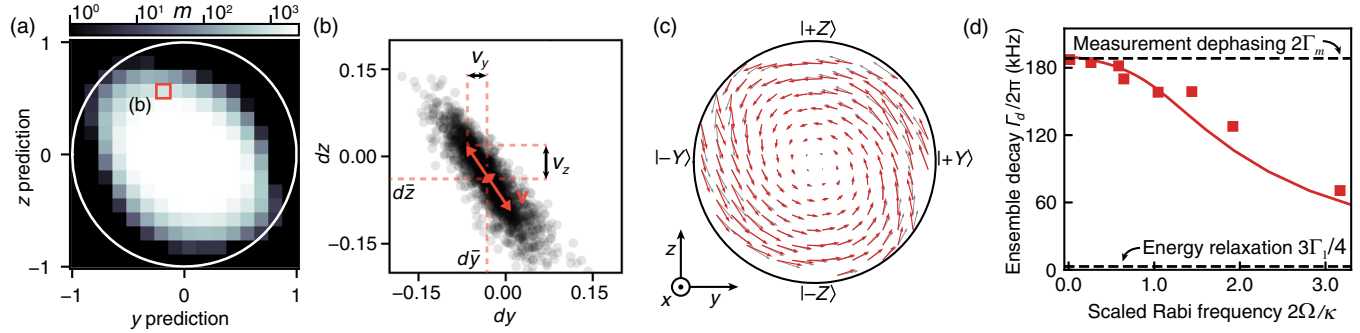

FIG. 3. LSTM trajectory processing and unraveling Lindbladian trajectory dynamics. (a) Histogram of the LSTM trajectories in the yz plane of the Bloch sphere for $2\Omega/\kappa = 0.6$. To extract trajectory drift and diffusion, we perform statistics on the Bloch vector increments $\{(dy_k, dz_k)\}_{k=0}^{m}$, where m is the number of samples in each bin; see color bar. (b) Statistics of the Bloch vector increments for the highlighted pixel in (a). The mean of the Bloch vector increments gives the average drift $(d\bar{y}, d\bar{z})$, whereas the scaled eigenvector of the covariance matrix ($\mathbf{v}$) gives the direction and magnitude of stochastic kicks due to measurement backaction. We compute $(d\bar{y}, d\bar{z})$ and $\mathbf{v}$ for each pixel in the yz plane to visualize the Lindbladian dynamics and measurement backaction. (c) The average drift $(d\bar{y}, d\bar{z})$ of trajectories binned in the yz plane ($2\Omega/\kappa = 0.6$, red arrows) reveals dynamics consistent with the applied Hamiltonian $H = \Omega\sigma_x/2$ and decay toward $y = 0$ at a rate Γ_d. The gray arrows represent a fit of the extracted pairs $(d\bar{y}, d\bar{z})$ to Eqs. (6) and (7), with Ω and Γ_d as free parameters. (d) The decay rate $\Gamma_d/2\pi$ (red squares) falls from the expected measurement dephasing rate $2\Gamma_m$ toward the bare qubit relaxation rate [45], as the qubit and resonator decouple. A master equation simulation of the qubit-resonator system (solid line) agrees qualitatively with the LSTM trajectories.

$$\sigma_{dy,dz} = \frac{1}{m}\sum_{k=0}^{m}(dy_k - d\bar{y})(dz_k - d\bar{z}), \tag{4}$$

$$\sigma_{dz}^2 = \frac{1}{m}\sum_{k=0}^{m}(dz_k - d\bar{z})^2, \tag{5}$$

is convenient to quantify diffusion since its largest eigenvalue $\lambda_{\max}$ represents the magnitude of the backaction, and the associated eigenvector $\boldsymbol{\xi}$ indicates the direction of the backaction. The vector $\mathbf{v} = (v_y, v_z) = \lambda_{\max}\boldsymbol{\xi}$ thus quantifies diffusion magnitude and direction, and is visualized in Fig. 3(b). By separating the drift and diffusion as described above, we can separately compare the average drift to Lindbladian dynamics, and the diffusion to the measurement backaction dynamics.

In the limit of a weak Rabi drive, the expected drift and diffusion can be directly calculated from the stochastic master equation in Eqs. (E1)–(E3). Therefore, in the plane of the Rabi drive ($x = 0$), we expect the average drift to take on the form

$$d\bar{y} = -\Gamma_d y dt + \Omega z dt, \tag{6}$$

$$d\bar{z} = -\Omega y dt. \tag{7}$$

Similarly, the expected diffusion is

$$\mathrm{Var}(dy) = -2\eta\Gamma_m z y dt, \tag{8}$$

$$\mathrm{Var}(dz) = 2\eta\Gamma_m(1 - z^2)dt. \tag{9}$$

Here, Γ_d is the ensemble average dephasing, Γ_m is the single quadrature measurement rate, and η is the total measurement efficiency. In the remainder of this section we first study the effect of the increasing Rabi drive on the average drift, and then study deviations from the backaction model of Eqs. (8) and (9).

In Fig. 3(c) we show an example of the average drift for $2\Omega/\kappa = 0.6$. The trajectory drift (red arrows) shows good agreement with the expected drift from the Hamiltonian $H = \Omega\sigma_x/2$ and dissipation pulling trajectories toward the origin at a rate $\Gamma_d/2\pi$ (gray arrows). From a fit of the average drift to Eqs. (6) and (7), using Ω and Γ_d as free parameters, we find $\Gamma_d/2\pi = 0.188 \pm 0.007$ MHz. For increasing Rabi drive strengths, Eqs. (6) and (7) continue to fit the average drifts. Interestingly, we observe a monotonic decrease in $\Gamma_d/2\pi$ [Fig. 3(d)], which we attribute to decoupling of the qubit from the weak measurement when the Rabi splitting of the dressed qubit exceeds the cavity linewidth [43]. This decoupling behavior is well captured by a master equation simulation of the qubit-resonator system [Fig. 3(d), solid line]. This master equation simulation also correctly shows a small Rabi drive asymptote $\Gamma_d \to 2\Gamma_m$ due to measurement dephasing [46] and another asymptote $\Gamma_d \to 3\Gamma_1/4$ [45] due to natural relaxation, when qubit and resonator are completely decoupled.

Since the average drift follows the expected Lindbladian dynamics even for fast Rabi rates, we now shift our attention to the measurement backaction, which we visualize by plotting the scaled eigenvectors $\mathbf{v}$ [Fig. 3(b)] of $C(y, z)$. At low Rabi frequencies ($2\Omega/\kappa = 0$), we find both informational backaction, vanishing at the poles $|\pm Z\rangle$ of the Bloch sphere, and phase backaction, consistent with a

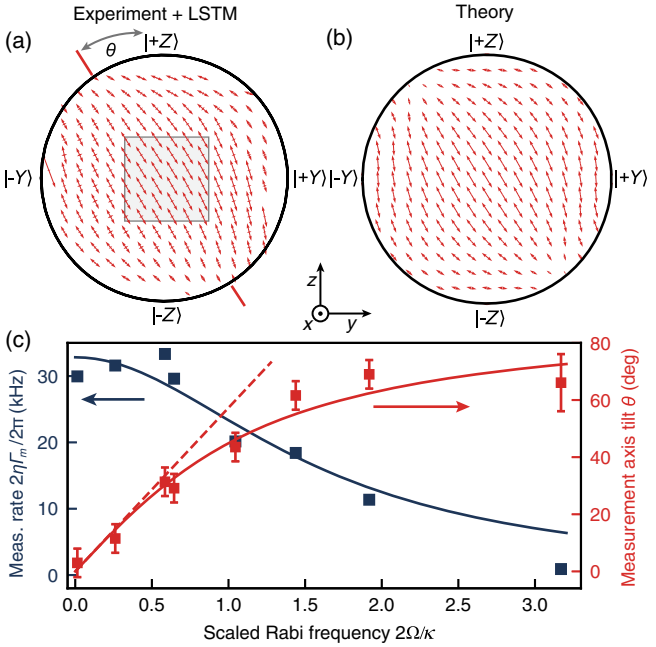

FIG. 4. Resonator memory corrections to the measurement backaction from LSTM trajectories. (a) Trajectory diffusion map for $2\Omega/\kappa = 0.6$, obtained by computing the eigenvectors of the covariance matrix Eq. (2) for trajectories binned in the yz plane. See also Figs. 3(a) and 3(b). The tilt toward $|\pm Y\rangle$ in the plane of the Rabi drive is in contrast to the prediction from Eq. (8) and (9). (b) The theory prediction, detailed in Appendix L, includes resonator memory effects and reproduces the tilt. (c) For small Ω/κ the measurement axis tilts linearly with the Rabi frequency $\theta = 2\Omega/\kappa$ (dashed line). The experimental tilt of the measurement eigenstates in the yz plane (red squares) is accompanied by a decrease in measurement rate (blue squares). Solid lines are fits to Eq. (11) (red) and Eq. (12) (blue), respectively. The error bars for θ are estimated from the imprecision in determining the tilt angle. The error bars for the measurement rate obtained from fitting are smaller than the markers.

dispersive heterodyne measurement (see Appendix H). Furthermore, a fit of the extracted diffusion (v_y, v_z) to Eqs. (8) and (9), along with $\Gamma_d/2\pi = 0.175$ MHz determined from Fig. 13(d), gives the efficiency of the measurement chain $\eta = 2\Gamma_m/\Gamma_d = 0.185 \pm 0.002$. This value is in agreement with an independent calibration measurement [$\eta = 0.188 \pm 0.003$ (see Fig. 13].

For large Rabi drives, we observe two significant corrections to the expected measurement backaction of Eqs. (8) and (9): a tilt of the measurement poles in the plane of the Rabi drive toward $|\pm Y\rangle$ and a reduced trajectory diffusion rate $2\eta\Gamma_m$ [Fig. 4(a)]. These corrections complicate extracting the diffusion rate. Therefore, we first determine the rotation θ from the average angle between $\mathbf{v}$ and the z axis from bins near the origin [Fig. 4(a), shaded square]. Next, we apply a coordinate rotation $(y', z') \mapsto (y\cos\theta - z\sin\theta, y\sin\theta + z\cos\theta)$ to undo the observed rotation, which now allows us to fit the diffusion $(v_{y'}, v_{z'})$ to Eqs. (8) and (9). The fit results for θ and the measurement rate $2\eta\Gamma_m$ are shown in Fig. 4(c), squares. Intuitively, the tilt occurs because the Rabi drive drags the qubit state counterclockwise in the yz plane while photons in the resonator measure the qubit z coordinate for a characteristic time τ_c. This simple picture results in a delay between qubit state and measurement record which is proportional to $\Omega\tau_c$, consistent with the linear increase in θ for small Rabi frequencies in Fig. 4(c).

To explain the effects of the slow measurement beyond small Ω/κ, we derive an analytical model for the resonator mode $\hat{a}$ in terms of the Bloch coordinate z, taking into account the resonator memory (see Appendix K). For a constant Rabi drive in the yz plane the weak measurement no longer acts purely along $\hat{z}$, but along

$$\hat{z}_{\text{eff}} = \int_0^t z(t-\tau)dP(\tau) = \sqrt{\eta_{\text{av}}}[\cos\theta\hat{z} - \sin\theta\hat{y}], \quad (10)$$

where P is the inverse Fourier transform of the Lorentzian resonator spectrum,

$$\theta = \arctan(\Omega\tau_c) \quad (11)$$

is the measurement axis tilt, $\tau_c = \int \tau P(\tau)d\tau = 2/\kappa$ is the predicted delay time, and η_{av} captures the reduction in measurement rate in the yz plane:

$$\Gamma_m(\Omega) = \eta_{\text{av}}\Gamma_m(0) = \frac{\Gamma_m(0)}{1 + (\Omega\tau_c)^2}. \quad (12)$$

From a simultaneous fit of both quantities in Fig. 4(c) to Eqs. (11) and (12), we find the resonator memory time $\tau_c = 0.20 \pm 0.01$ μs, which agrees with the relaxation time $2/\kappa = 0.2$ μs from spectroscopic measurements. For Rabi frequencies exceeding $2\Omega/\kappa \approx 2.0$ the rate of trajectory diffusion in the yz plane is greatly suppressed, showing that a fast Rabi frequency protects the qubit from measurement backaction [43].

The results of Fig. 3 and 4 demonstrate accurate estimation of various decay rates, measurement efficiency, and the memory time of the system from a simple trajectory decomposition that requires no prior knowledge of the resonator memory. This makes the LSTM a useful tool in the context of parameter estimation from weak measurements [40,47]. In addition, the corrections to the measurement backaction shed further light on the breakdown of Baysian filters for trajectory reconstruction.

V. UNCOVERING HIDDEN TIME DEPENDENCIES

A major advantage of the LSTM is that trajectory reconstruction does not require prior knowledge of system and environment parameters, which may fluctuate or may

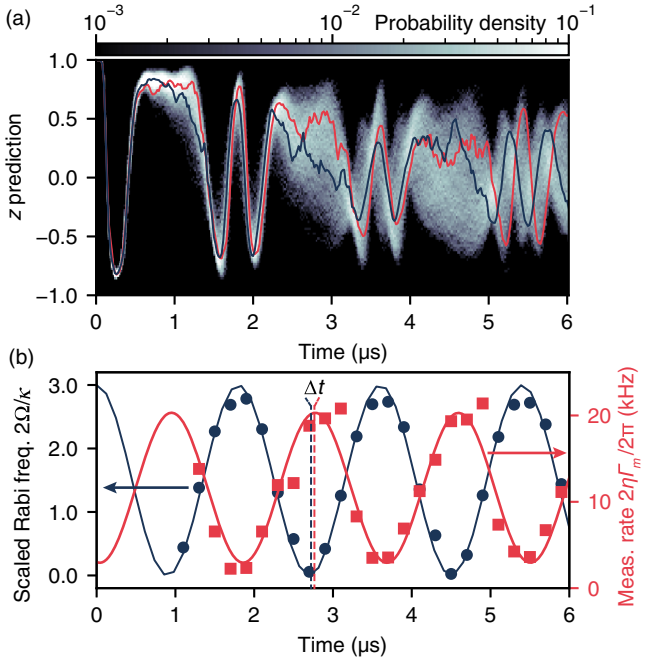

FIG. 5. Uncovering hidden time dependencies from trajectories subject to sinusoidal Rabi drive modulation. (a) A histogram of the reconstructed z coordinate shows periods of increased and decreased diffusion, also visible in two sample trajectories (red and blue). (b) Repeated analysis of the LSTM trajectories for successive 0.2-μs-long time windows. For each window we extract the instantaneous Rabi frequency and measurement rate by fitting the average drift to Eqs. (6) and (7) and diffusion to Eqs. (8) and (9), respectively. The extracted measurement rate $2\eta\Gamma_m/2\pi$ (red squares) shows maxima during periods of large diffusion in (a). The solid red line is a simple sinusoidal fit to the red squares. The instantaneous Rabi frequencies (blue dots), obtained from fitting trajectories to the Lindbladian part of the reduced qubit master equation, are consistent with the applied drive amplitude (solid blue line). Fit results for $t < 1$ μs are not shown, since early during the evolution trajectories have not spread sufficiently for accurate fit results.

be hard to calibrate. To highlight this advantage, we perform a new set of weak measurements where we vary the Rabi frequency sinusoidally with a period of 1.8 μs, thereby obscuring the prior information necessary for traditional trajectory reconstruction methods. The ability to reconstruct trajectories with *a priori* unknown time dependencies is important for a wide range of experiments, such as continuous quantum error correction or weak measurements of a highly non-Markovian environment [48].

The resulting histogram of LSTM trajectories [Fig. 5(a)] shows alternating periods of trajectory bunching, where the trajectories are protected from measurement backaction, and diffusion, where the weak measurement clearly imparts stochastic kicks to the qubit. We unravel this nontrivial dynamics by dividing the LSTM trajectories in 0.2-μs-long time windows and analyze the average drift and diffusion as described in the previous section. During each 0.2 μs window, the trajectories only partly cover the yz plane, but we can still extract the Rabi frequency from the average drift and measurement rate from the diffusion [Fig. 5(b)]. Without prior knowledge of the time dependence of Ω, Γ_d, or Γ_m, we correctly recover the sinusoidal applied Rabi frequency [Fig. 5(b), blue], and find that it anticorrelates strongly with the extracted measurement rate [Fig. 5(b), red], consistent with the results in Fig. 4(c). A sinusoidal fit to the instantaneous measurement rate shows a delay of $\Delta t = 0.14 \pm 0.06$ μs with respect to the instantaneous Rabi frequency, consistent with the resonator relaxation time $2/\kappa$. The results of Fig. 5 demonstrate that the LSTM correctly reconstructs trajectories with hidden time dependencies, and that parameters can be extracted with high time resolution, ultimately limited by the time step dt of the voltage record.

VI. CONCLUSIONS

We have shown that LSTM recurrent neural networks can outperform standard quantum filters if the qubit dynamics is faster than the resonator linewidth. In this previously unexplored regime of weak measurement, the LSTM correctly recovers Lindbladian dynamics and distorted measurement backaction without prior knowledge of the resonator memory. For the strongest Rabi drives the LSTM validation error remains low; however, in this case the resonator temporally averages most of the qubit signal, reducing the best trajectory estimate to an ensemble-averaged trajectory. Therefore, the LSTM is most useful in the intermediate regime $2\Omega/\kappa < 2$, where the measurement backaction is distorted, but the backaction strength is nonzero and trajectories are therefore nontrivial. We further showed that the LSTM is not constrained to constant Hamiltonians. Even when underlying parameters vary on a submicrosecond timescale the LSTM accurately recovers quantum trajectories, which allows parameter estimation with high time resolution, and may enable trajectory reconstruction for qubits connected to strong non-Markovian environments [38,49] where coupling rates can vary with time.

To further improve the LSTM accuracy for applications such as parameter estimation, it is possible to add physical constraints to the LSTM loss function [40,50,51]. In addition, other neural network architectures such as tensor networks or models based on dilated causal convolutions [52] may improve prediction accuracy, though such comparisons require more training data. Finally, a possible extension of our LSTM is state tracking in larger Hilbert spaces. This requires more projective measurement pulses (e.g., at least 4 for single qutrits or 9 for two qubits) compared with the single qubit state tracking, which only requires 3 tomography axes. Given the single qubit data acquisition time of approximately 3 h, extending our LSTM tracking to single qutrits or two qubits seems feasible,

though optimizing data transfer from the digitizer to computer memory may be desirable to keep data acquisition efficient.

The code and data used for training are available on Github [53].

ACKNOWLEDGMENTS

The authors would like to thank K. Nowrouzi, B. Mitchell, and A. Morvan for experimental assistance, and L. B. Nguyen and Y. Kim for valuable comments on the manuscript. We gratefully acknowledge insightful conversations with L. Martin, É. Genois, A. DiPaolo, J. Gross, and M. Laskin. This work was supported by the U.S. Army Research Laboratory and the U.S. Army Research Office under Contract Grant No. W911NF-17-S-0008. L. B., S. G. and J. D. acknowledge additional support from NSF-BSF Grant Award No. 1915015.

APPENDIX A: EXPERIMENTAL SETUP

The experiments in this work were performed with a chip of 8 superconducting transmon qubits, cooled to 10 mK in a BlueFors XLD dilution refrigerator. Room-temperature and cryogenic electronics for qubit control and measurement are shown in Fig. 6. Qubit control pulses are generated by up-conversion of intermediate frequency (IF) pulses generated by a Keysight PXI arbitrary waveform generator (AWG) via in-phase quadrature (IQ) modulation of a continuous wave (cw) local oscillator (LO) tone, sourced by a Keysight MXG N5183B at 5.415 GHz. Both I and Q components are sourced at 1 GS/s between 66 and 261 MHz. The phase and dc offsets between the I and Q waveforms are tuned to eliminate the opposite sideband and LO leakage due to mixer nonidealities, while bandpass filtering at room temperature reduces noise from the AWG. Readout pulses are generated with the same AWG and are similarly up-converted with a 6.83 GHz LO tone from a

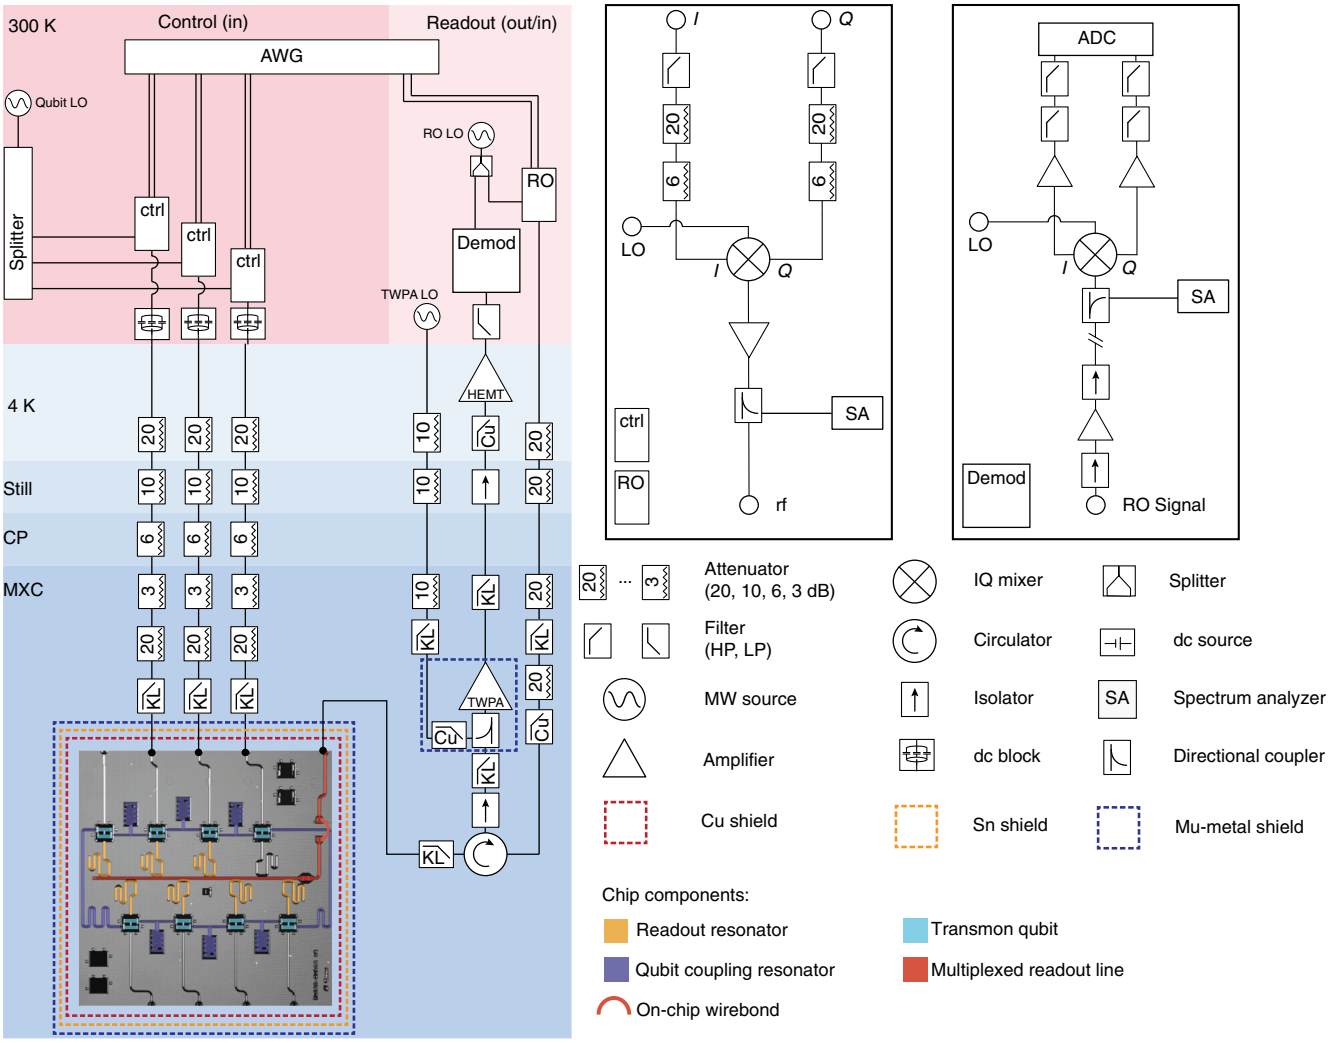

FIG. 6. Microwave electronics inside and outside the dilution refrigerator. CP, cold plate (100 mK); MXC, mixing chamber (10 mK). Any additional description of terms can be found in Appendix A. Abbreviations: ADC, analog to digital converter; MW, microwave; KL, name of a commercial microwave filter brand.

separate Keysight MXG N5183B. Control pulses are passed through a dc block, and all pulses traveling to the sample are attenuated at each stage, with a further K&L low-pass filter at the base stage. The up-conversion chain schematic for both the qubit control and readout lines is detailed in the inset labeled "ctrl" and "RO."

Reflected measurement signals are redirected by a circulator to a measurement chain outfitted with superconducting coaxial cable, where they are amplified by a traveling wave parametric amplifier at 10 mK, a high electron mobility transistor (HEMT) at 4 K, and a low-noise room-temperature amplifier, before being down-coverted to IQ IF components with the same 6.83 GHz LO tone used for RO up-conversion. The down-converted signals are then amplified and filtered to reduce high-frequency amplifier noise before being digitized at 1 GS/s by an Alazar analog-to-digital converter (ADC) and demodulated in software. The down-conversion chain schematic is detailed in the inset labeled "Demod."

The TWPA is pumped with a cw tone, sourced by a Hittite HMC M2100 at 7.42 GHz and 16.5 dBm before being attenuated and filtered by low-pass K&L and copper powder filters.

APPENDIX B: DEVICE PARAMETERS AND STABILITY

Table I summarizes the relevant parameters of the qubit and readout resonator used in this work. We obtain the resonator frequency ω_{res} and cavity linewidth $\kappa/2\pi$ through spectroscopic measurements, and we measure the qubit frequency ω_{ge}, the anharmonicity $\alpha/2\pi$, and $T_{2,\mathrm{Ramsey}}$ through repeated Ramsey measurements on the g-e transition (ω_{ge}, $T_{2,\mathrm{Ramsey}}$) and e-f transition ($\alpha/2\pi$). Lastly, we fit the weak measurement strength $\varepsilon/2\pi$ by comparing the measured Rabi decay time as function of Rabi frequency to a master equation simulation using QuTiP [54]. The error bar on ε is due to uncertainty in the fit.

TABLE I. Calibrated resonator and qubit parameters for the device used in this work. Note that ω_{res} and ω_{ge} are the bare readout resonator and qubit frequencies, respectively, and that T_1 and $T_{2,\mathrm{Ramsey}}$ were measured without a weak measurement tone ($\varepsilon/2\pi = 0$). The error bars on T_1 and $T_{2,\mathrm{Ramsey}}$ reflect the standard deviation of 100 repeated measurements during a 3 h window.

Parameter	Description	Measured value
$\omega_{\mathrm{res}}/2\pi$	Bare cavity frequency	6679 MHz
$\kappa/2\pi$	Cavity linewidth	1.56 MHz
$\omega_{ge}/2\pi$	Transmon frequency (empty cavity)	5473 MHz
$\alpha/2\pi$	Transmon anharmonicity	−270 MHz
T_1	Transmon relaxation time ($\Omega = 0$)	61 ± 7 μs
$T_{2,\mathrm{Ramsey}}$	Ramsey dephasing time (empty cavity)	70 ± 9 μs
$\chi/2\pi$	Half of the dispersive cavity shift	0.47 MHz
$\varepsilon/2\pi$	Weak measurement strength	0.43 ± 0.01 MHz

Acquiring a training dataset (i.e., one initial qubit state with three tomography axes) for a single Rabi frequency takes approximately 3 h. As the neural network learns a representation of the qubit Hamiltonian, it is important that the parameters remain stable during the acquisition. In Sec. III we have shown that the accuracy of the LSTM is currently limited by the dataset size. However, as the cross-entropy loss saturates for larger datasets, we expect experimental instabilities to play a bigger role in limiting LSTM accuracy. Therefore, we briefly discuss two known sources of instabilities: coherence fluctuations and readout fidelity fluctuations.

Since readout infidelity and unexpected decay or dephasing directly affect the training labels, we have characterized the stability of T_1, $T_{2,\mathrm{Ramsey}}$ as well as the readout fidelity over a 3 h time span. For T_1 and $T_{2,\mathrm{Ramsey}}$ we observe fluctuations on the order of 10%–15% of their respective values, which is likely due to nearby two-level systems that couple weakly to the qubit. Since the fluctuations are modest, and average values of 61 and 70 μs remain much larger than the maximum trajectory length of 8.0 μs, we do not expect decay time fluctuations to affect the current performance of the neural network.

We characterize readout infidelity by repeatedly preparing the qubit in the ground and excited state followed by a measurement of the qubit state. We only observe small fluctuations of the correctly classified state over a 3 h window and find average correctly classified probabilities $P(0|0) = 0.995 \pm 0.005$ and $P(1|1) = 0.984 \pm 0.004$. We expect that the average readout infidelity of $1 - \frac{1}{2}[P(0|0) + P(1|1)] = 1.0\%$ may introduce a small bias to the trained neural network, but fluctuations in this quantity do not play a significant role.

The qubit studied in this work is part of a larger 8-qubit quantum processor, and is capacitively coupled to two neighboring qubits, each via a coupling bus with an effective coupling strength of approximately 2 MHz. To avoid any undesired qubit dynamics due to the neighboring qubits, we postselect the voltage records, conditioning on all qubits starting in the ground state before the Rabi sequence, and the neighboring qubits remaining in the ground state after the Rabi sequence. Typically the post-selection removes 5% of the voltage records.

APPENDIX C: PROCESSING OF MEASUREMENT RECORDS

Traditional processing of weak measurement records includes digital filtering and choosing a coarse-grained dt that minimizes temporal correlations in the measurement record, to ensure that the appropriate stochastic master equation can be applied faithfully to the measured data [16,17]. As explained in the main text, for fast qubit dynamics it becomes impossible to coarse grain the record

TABLE II. Properties of the datasets used for Figs. 2–4 in the main text. The column header "Initial qubit state" refers to the prepared initial qubit state. "Dataset size" is the total number of trajectories available for each dataset, which includes a fraction of trajectories used for analysis ("Trajectories for LSTM analysis"). The subsequent columns describe the longest Rabi drive duration (T_m), the sampling increment of the Rabi drive duration, and the coarse graining of the voltage records. Finally, "Sequence length" refers to the number of samples in each voltage record, which is given by the ratio of the longest T_m and dt. Note that all voltage records within a dataset have equal sequence length, since sequences with Rabi drive durations smaller than T_m are zero padded.

$2\Omega/\kappa$	Initial qubit state	Dataset size	Trajectories for LSTM analysis	Longest T_m (μs)	T_m increment (ns)	dt (ns)	Sequence length
0.0	$\|+Y\rangle$	454×10^3	46 579	8.0	200	40	200
0.0	$\|-Y\rangle$	454×10^3	46 540	8.0	200	40	200
0.2	$\|+Z\rangle$	283×10^3	29 027	8.0	200	40	200
0.6	$\|+Z\rangle$	283×10^3	29 029	8.0	200	40	200
1.0	$\|+Z\rangle$	283×10^3	29 014	8.0	200	40	200
1.4	$\|+Z\rangle$	283×10^3	29 062	8.0	200	40	200
2.0	$\|+Z\rangle$	353×10^3	36 054	5.0	100	20	250
3.2	$\|+Z\rangle$	354×10^3	36 118	5.0	100	20	250

in a way that produces uncorrelated measurement records without undersampling the data. We strike a balance between undersampling and correlation by coarse graining the measurement records as listed in Table II. Note that all records are sampled faster than the uncorrelated condition $dt > 2/\kappa = 0.2$ μs, and the two fastest datasets are sampled more finely than all others ($dt = 20$ ns). Feeding the Bayesian quantum filter correlated measurement records is expected to contribute to the rms error in Fig. 2(d). Interestingly, the Bayesian filter remains surprisingly accurate up to $2\Omega/\kappa \approx 1$.

APPENDIX D: TRAINING OF THE NEURAL NETWORK

1. Training data acquisition details

The pulse sequence for training data acquisition is given in Fig. 7(a). For each experimental run, the qubit is heralded in the ground state $|0\rangle$ with a strong measurement, resulting in a large mean photon occupation number $\bar{n}$ in the cavity. The measurement tone amplitude is then reduced to the weak measurement setpoint ($\bar{n} \approx 0.3$), and a 5 μs idling time allows $\bar{n}$ to reach steady state. Next, we prepare the qubit in one of the cardinal points of the Bloch sphere with a 30 ns preparation pulse, which is detuned from ω_{ge} to account for the Stark shift [55] induced by the nonzero cavity $\bar{n}$. At $t = 0$ we introduce σ_x dynamics by quickly ramping up the Rabi drive (ramp time 5 ns), while continuing to weakly measure the qubit along σ_z. This yields voltage records (I_n, Q_n) of variable duration T_m, which ranges between 0 and 8 μs. Finally, a projective measurement along one of the principal axes of the Bloch sphere $\sigma_{x,y,z}$ is performed by first applying a single qubit gate to rotate the qubit and then a strong measurement ($\bar{n} \gg 1$). The projective measurement results are binary variables $P_i \in \{0, 1\}$ and represent the training labels for the neural network.

2. Training details

In this work we use a long short-term memory recurrent neural network, a model that recurrently updates a layer of virtual neuronlike nodes as it processes a time-serialized input, and features a memory cell to retain past information of the time series [56,57]. At time t, the layer is encoded in a vector $\vec{h}_t$, and propagates to time $t + 1$.

A typical dataset with single initial qubit state, and a single Rabi frequency, consists of $N \approx 4 \times 10^3$ repetitions of $M = 40$ different length voltage records, repeated for three tomography axes. We use 90% of the approximately $3MN = 4 \times 10^5$ voltage records for training and the remaining 10% for cross validation (see Table II).

The training code uses the Tensorflow 2 library in PYTHON and consists of several epochs where training samples are fed forward and backward through the neural network in minibatches of size $N_b = 512$. During the training weights and biases of the network are adjusted to minimize a cost function $\mathcal{L}$, which contains the following three components (similar to Ref. [40]).

(i) A cross-entropy loss associated with the projective measurement result at the end of a voltage record,

$$\mathcal{L}_{\mathrm{CE}} = -\frac{1}{N_b}\sum_{i=1}^{N_b} P_i \log(s_i) + (1 - P_i)\log(1 - s_i),$$

where P_i and s_i are the tomography result (0 or 1) and probability assigned by neural network ($0 \le s_i \le 1$) for voltage record i in the minibatch, respectively.

(ii) A mean-squared error for deviating from the known initial state at time $t = 0$, $\mathcal{L}_{\mathrm{init}}$.

(iii) A rectified linear unit activation on the purity of the quantum state $\mathcal{L}_{\mathrm{purity}}$ at all times, which tries to enforce that the Bloch vector does not lie outside the Bloch sphere for all t.

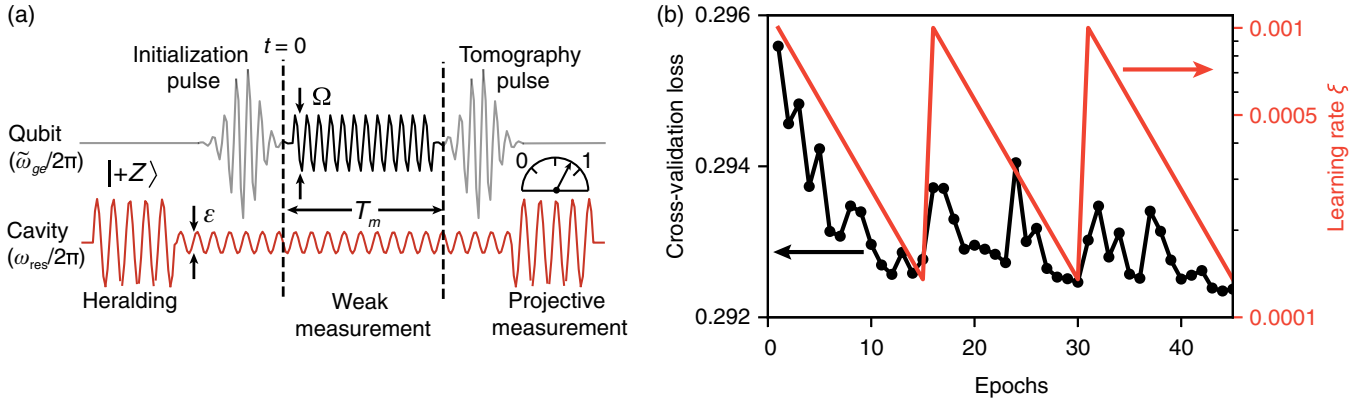

FIG. 7. (a) Experimental pulse sequence for training data acquisition. Time axis is not to scale. (b) Cost function (black, left-hand axis) computed on the cross-validation dataset during a typical training session. Temporary increases in the cost function are expected as the learning rate (red, right-hand axis) increases periodically.

We construct the loss function as a weighted sum of the components above, but note that the main contribution to the loss function comes from $\mathcal{L}_{\mathrm{CE}}$.

We use the ADAM optimizer [58] to obtain the LSTM parameters that minimize the cost function $\mathcal{L}$, and adjust the step size of the stochastic gradient descent (i.e., learning rate) ξ in a cyclical fashion [Fig. 7(b)] to prevent getting stuck in local minima [59]. We have experimented with adding dropout and regularization, techniques which aim to improve generalization on unseen data [60], but have not observed improvements in $\mathcal{L}$ computed on the validation dataset. Since dropout randomly resets a fraction of trainable parameters between epochs, it destroys temporal correlations between successive LSTM cells which may be hard to learn from the training data. In addition, we believe that the lack of improvement could be due to the noisy nature of the voltage records and to the stochastic nature of the projective measurements (training labels), which could prevent overfitting even in the absence of dropout and regularization.

During the training process we monitor the value of $\mathcal{L}$ for both the training and cross-validation samples [for an example, see Fig. 7(b)]. The neural network is converged when $\mathcal{L}$ computed on the training dataset no longer decreases after an annealing cycle [33]. The cross-validation samples are then passed through the network once more and we use the resulting trajectories in Figs. 2–5 of the main text. For our largest dataset of 0.8×10^5 voltage records, the entire training process completes within approximately 20 min thanks to the processing power of the graphics processing unit (NVidia GeForce RTX 2080 Ti).

APPENDIX E: TRAJECTORY HISTOGRAMS

The trajectory histograms shown in Fig. 2(a) of the main text show that for $t \gg \Gamma_d^{-1}$ trajectories never cross a threshold in z, and therefore seem constrained to a reduced-radius Bloch sphere. In this appendix, we derive this effect from the stochastic master equation.

We consider the stochastic master equation (in Itô form) for heterodyne measurement of a driven qubit, where the Rabi drive is applied in the yz plane [61]:

$$dx = -\Gamma_d x dt - \sqrt{2\eta\Gamma_m} z x dW_1 - \sqrt{2\eta\Gamma_m} y dW_2, \quad \text{(E1)}$$

$$dy = -\Gamma_d y dt + \Omega z dt - \sqrt{2\eta\Gamma_m} z y dW_1 + \sqrt{2\eta\Gamma_m} x dW_2, \quad \text{(E2)}$$

$$dz = -\Omega y dt + \sqrt{2\eta\Gamma_m}(1 - z^2) dW_1. \quad \text{(E3)}$$

These equations include informational backaction (dW_1) causing trajectories to collapse to the poles $|\pm Z\rangle$ at a rate $\eta\Gamma_m$, phase backaction (dW_2) at an equal rate due to the quantum fluctuations in the cavity photon number $\bar{n}$, and dephasing toward the origin of the Bloch sphere at a rate $2(1-\eta)\Gamma_m$. Note that we define Γ_m as the decoherence rate due to measurement of a single quadrature, and since we measure both I and Q quadratures, the total dephasing rate includes a factor of 2. Further note that for a unit efficiency measurement setup, individual trajectories diffuse quickly but remain pure. Therefore, the reduced radius is due to limited measurement efficiency η.

To further quantify the reduced radius observed in the experimental trajectories, we calculate the mean change of the purity, defined as $|\mathbf{r}|^2 = x^2 + y^2 + z^2$:

$$\left\langle \frac{d|\mathbf{r}|^2}{dt} \right\rangle = -2[(1-\eta)\Gamma_m + \Gamma_m](|\mathbf{r}|^2 - z^2) + 2\eta\Gamma_m(1 - 2z^2) + 2\eta\Gamma_m|\mathbf{r}|^2 z^2. \quad \text{(E4)}$$

We can simplify this expression in the fast Rabi regime ($\Omega \gg \Gamma_m$), where $z^2 \approx |\mathbf{r}|^2/2$ evaluated over one Rabi cycle. This yields

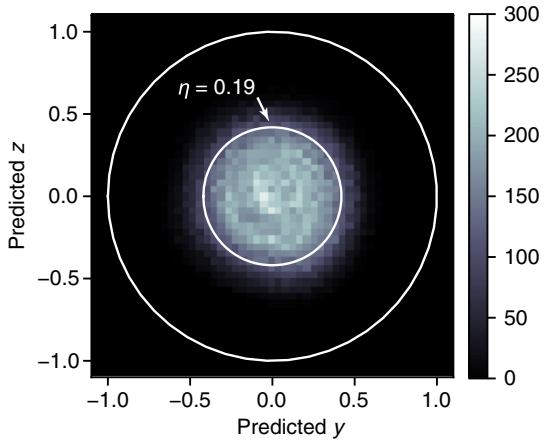

FIG. 8. Trajectory histogram for a fast Rabi drive $2\Omega/\kappa = 2.0$, postselected for steady state (times longer than Γ_d^{-1}). The inner circle is a prediction for the purity of individual trajectories from Eq. (E6) with η calibrated from Appendix J.

$$\left\langle \frac{d|\mathbf{r}|^2}{dt} \right\rangle = -2\eta\Gamma_m - (2+\eta)\Gamma_m|\mathbf{r}|^2 + \eta\Gamma_m|\mathbf{r}|^4. \tag{E5}$$

From Eq. (E5) we calculate the steady-state solution for the reduced radius,

$$|\mathbf{r}| = \sqrt{\frac{2+\eta}{2\eta} - \sqrt{\left(\frac{2+\eta}{2\eta}\right)^2 - 2}}, \tag{E6}$$

which depends entirely on the measurement efficiency, and requires a small correction if the environmental dephasing rate is nonzero. Using the efficiency calibration from Appendix J, we find that the predicted radius $|\mathbf{r}|$ agrees well with the observed trajectories histogram shown in Fig. 8 for $2\Omega/\kappa = 2.0$. While the same physics applies to the histogram in Fig. 2 of the main text, the predicted $|\mathbf{r}|$ is inaccurate for this moderate Rabi rate $\Omega/\Gamma_m \approx 1$, since the approximation $\Omega \gg \Gamma_m$ is invalid. In this regime further corrections to Eq. (E5) are necessary, but these are beyond the scope of this work.

APPENDIX F: ACCURACY OF THE PREDICTED TRAJECTORIES BY THE LSTM

In Fig. 2(b) of the main text, we demonstrate the accuracy of the LSTM trajectories for a slow Rabi drive $2\Omega/\kappa = 0.2$. The purpose of this appendix is to define the accuracy and to discuss error sources.

To quantify the LSTM accuracy, we bin all predictions immediately before a projective measurement associated with each trajectory and average the projective measurement results for trajectories with similar predictions. This technique is frequently used to experimentally verify an estimate of the qubit state with an independent measurement [15,62]. We use the following definition of the weighted root-mean-square error for the x coordinate (y and z are similar):

$$\varepsilon_x = \sqrt{\sum_{i=1}^{N_x} N_i \Big(\underbrace{x_{\mathrm{LSTM},i}}_{\text{LSTM prediction}} - \underbrace{\langle x_{\mathrm{proj},i}\rangle}_{\text{av tomography}} \Big)^2 \Big/ \sum_{i=1}^{N_x} N_i}, \tag{F1}$$

where N_x is the number of bins, $x_{\mathrm{LSTM},i} \in [-1,1]$ is the predicted value by the LSTM for bin i, and $\langle x_{\mathrm{proj},i}\rangle$ is the average of projective measurements for predictions within $\delta = 2/N_x$ of $x_{\mathrm{LSTM},i}$. The definition of Eq. (F1) captures any inaccuracy caused by the LSTM (e.g., inefficient model, insufficient training data) but is also sensitive to the uncertainty in our best estimate of the ground truth $\langle x_{\mathrm{proj},i}\rangle$, which scales as $1/\sqrt{N_i}$. This becomes a problem for rare predictions, when the number of trajectories N_i is small. Therefore, by weighing each bin by the number of predictions in that bin N_i, ε_x becomes less sensitive to limited statistics for rare predictions.

In Fig. 9 we show that the uncertainty in the projective measurement dominates any errors caused by the LSTM, since the absolute error $|(x_{\mathrm{LSTM},i} - \langle x_{\mathrm{proj},i}\rangle)|$ closely

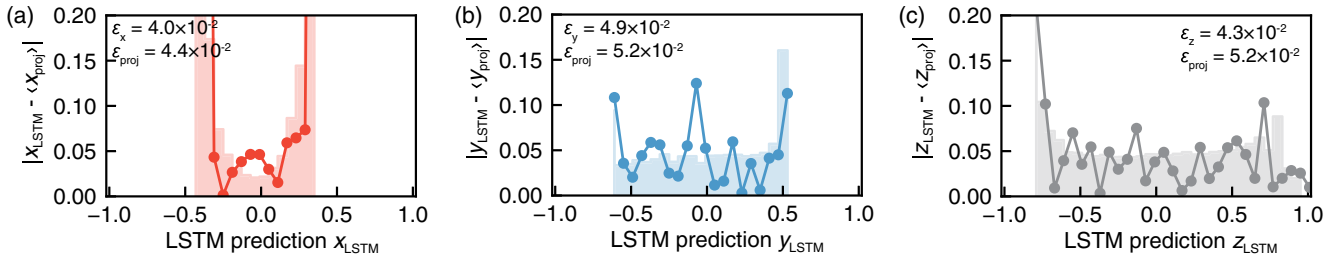

FIG. 9. Comparison of the LSTM prediction error (dots) and expected uncertainty in the estimate of the ground truth $\langle x_{\mathrm{proj},i}\rangle$ [histograms calculated from Eq. (F2)] for Bloch coordinates x (a), y (b), and z (c). $\varepsilon_{x,y,z}$ are the weighted rms errors calculated from LSTM predictions, and $\varepsilon_{\mathrm{proj}}$ is calculated from Eq. (F1) while setting $x_{\mathrm{LSTM},i} - \langle x_{\mathrm{proj},i}\rangle = \sigma_{\langle x_{\mathrm{proj},i}\rangle}$. $\varepsilon_{\mathrm{proj}}$ thus represents the weighted rms error in case of perfect LSTM predictions and 1σ statistical fluctuation in $\langle x_{\mathrm{proj},i}\rangle$ due to a limited number of trajectories. Since both quantities follow similar trends and their magnitudes $\varepsilon_{x,y,z}$ and $\varepsilon_{\mathrm{proj}}$ are consistent, we conclude that our accuracy measure is limited by the number of trajectories.

follows the expected uncertainty in $\langle x_{\mathrm{proj},i}\rangle$, estimated from the Bernoulli distribution:

$$\sigma_{\langle x_{\mathrm{proj},i}\rangle} = \sqrt{\frac{1-\langle x_{\mathrm{proj},i}\rangle^2}{N_i}}. \tag{F2}$$

We observe similar trends in the absolute error and the uncertainty estimated from Eq. (F2) for all Bloch coordinates [Figs. 9(a)–9(c)], and error magnitudes are consistent. In other words, our ability to quantify the LSTM error is limited by our knowledge of the true quantum state. To further assess the LSTM performance, more trajectories are needed to reduce our uncertainty in the estimate of the true quantum state. Alternatively, it is possible to train the LSTM on simulated trajectories for which the true quantum state is known [40].

APPENDIX G: CAVITY CORRECTIONS TO THE BAYESIAN FILTER

Quantum trajectory theory describes how an observer's estimate of a quantum state evolves as it is updated with a weak measurement record [8,62]. The measurement record is translated to quantum state evolution by first applying the unitary evolution and then updating the state with the measurement record at each time step, as the backaction of the measurement on the state can be derived [25]. Thus, with knowledge of the initial state and the Hamiltonian driving unitary evolution, the density matrix $\rho_{ij}(t_i)$ can be repeatedly updated.

We follow the procedure outlined in Refs. [10,63]. For each time t_i in the voltage record, we first apply the Lindbladian dynamics,

$$\rho'(t_i) = \exp(i\lambda dt)\rho(t_i), \tag{G1}$$

where λ is a 4×4 matrix and $\rho(t_i) = (\rho_{00}, \rho_{01}, \rho_{10}, \rho_{11})$. Including a Rabi drive detuning Δ and a Rabi drive misalignment Ω_y, λ takes the form

$$\lambda = \begin{pmatrix} 0 & \frac{\Omega_+}{2} & -\frac{\Omega_-}{2} & -i\Gamma_1 \\ \frac{\Omega_-}{2} & -\Delta & 0 & -\frac{\Omega_-}{2} \\ -\frac{\Omega_+}{2} & 0 & \Delta & \frac{\Omega_+}{2} \\ 0 & -\frac{\Omega_+}{2} & \frac{\Omega_-}{2} & i\Gamma_1 \end{pmatrix}, \tag{G2}$$

$$\Omega_\pm = \Omega_x \pm i\Omega_y, \tag{G3}$$

$$\Delta = \omega_{ge} - \omega_{\mathrm{Rabi}} \approx 0, \tag{G4}$$

$$\Gamma_1 = \frac{1}{T_1}. \tag{G5}$$

Note that λ describes the same Lindbladian dynamics that leads to Eqs. (E1)–(E3) provided $\Omega_y = 0$, and $dt = t_{i+1} - t_i$ is the time increment of the voltage record. The parameters $\Omega_{x,y}$, Δ, and Γ_1 are obtained from a fit to the projective measurement results. Next, we update the density matrix taking into account measurement record $[I(t_i), Q(t_i)]$ and the partially updated density matrix $\rho'(t_i)$:

$$\rho_{11}(t_{i+1}) = \frac{\rho'_{11}(t_i)/\rho'_{00}(t_i)e^{\alpha_i}}{1+\rho'_{11}(t_i)/\rho'_{00}(t_i)e^{\alpha_i}}, \tag{G6}$$

$$\rho_{00}(t_{i+1}) = 1 - \rho_{11}(t_{i+1}), \tag{G7}$$

$$\rho_{10}(t_{i+1}) = \rho'_{10}(t_i)\sqrt{\frac{\rho_{11}(t_{i+1})\rho_{00}(t_{i+1})}{\rho'_{11}(t_i)\rho'_{00}(t_i)}}e^{-i\beta_i}e^{-2(1-\eta)\Gamma_m dt}, \tag{G8}$$

$$\rho_{01}(t_{i+1}) = \rho^\dagger_{10}(t_{i+1}), \tag{G9}$$

where $\alpha_i = \tilde{I}_i\Delta I/\sigma^2$, $\beta_i = \tilde{Q}_i\Delta I/2\sigma^2$, σ^2 is the variance of the noise in the voltage record, and finally,

$$\tilde{I}_i = I(t_i) - \frac{(I_0+I_1)}{2}, \tag{G10}$$

$$\tilde{Q}_i = Q(t_i) - Q_0, \tag{G11}$$

$$\Delta I = I_1 - I_0. \tag{G12}$$

I_0 and I_1 are the steady-state coherent-state amplitudes conditioned on the qubit in the ground state and excited state, respectively. These state update equations are exact in the absence of a Rabi drive, and if the qubit decay rates are small compared with κ [10]. Figure 2(d) in the main text shows that these equations accurately reconstruct trajectories up to $2\Omega/\kappa \approx 1$ even though they do not contain the cavity.

While ΔI can be obtained with a simple calibration measurement when $\Omega = 0$, it is impossible to calibrate ΔI in the presence of a Rabi drive, since the measurement term (σ_z) and Rabi drive (σ_x) do not commute. One might assume that ΔI does not change with the applied Rabi frequency, but this leads to a relatively large rms error even at moderate Rabi drives. We can improve the performance of the Bayesian filter reconstruction at moderate Rabi drive by adjusting ΔI based on prior information obtained from a QuTiP numerical simulation which includes qubit and coupling to a cavity. From these simulations [Fig. 10(a)] we find that the steady-state coherent-state amplitudes conditioned on the qubit state ($\alpha_{g,e}$) substantially decrease as the Rabi frequency increases. Therefore, to correct for cavity effects, we scale ΔI by the steady-state value of $\alpha_e(\Omega) - \alpha_g(\Omega)/[\alpha_e(0) - \alpha_g(0)]$ [Fig. 10(b)], which yields the triangles in Fig. 2(d) of the main text with substantially smaller rms error.

While the scaling factor for ΔI partially corrects for the cavity memory effect, the rms error remains large for

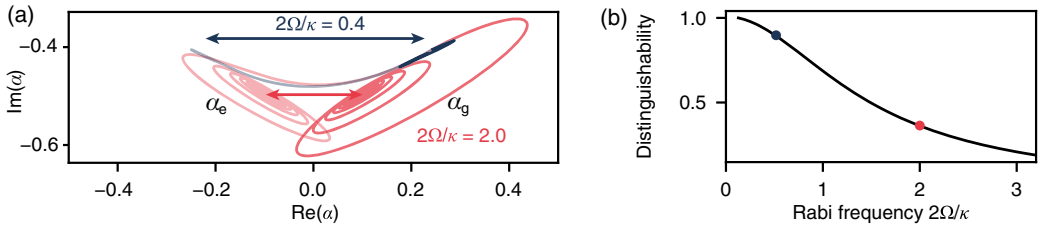

FIG. 10. (a) QuTiP master equation simulation of a driven qubit dispersively coupled to the readout cavity, showing coherent state amplitudes versus time, conditioned on the qubit state. We show two pairs of α_g and α_e for $2\Omega/\kappa = 0.4$ and $2\Omega/\kappa = 2.0$ obtained from the expectation value $\alpha_g \approx \mathrm{Tr}(\rho_{q+c} a|g\rangle\langle g|/P_g)$, where a is the cavity photon annihilation operator and ρ_{q+c} is the joint qubit-cavity density matrix. (b) We use the relative distinguishability, defined as $\alpha_g(\Omega) - \alpha_e(\Omega)/\alpha_g(0) - \alpha_e(0)$, to scale ΔI and partially correct for cavity effects in the Bayesian filter method.

$2\Omega/\kappa \gg 1$, since temporal correlations in the voltage record are large. In this case the simple steady-state Bayesian state update equations are no longer valid, and the most straightforward way to further reduce the validation error is by including the full cavity in the stochastic master equation. However, as stated in the main text, this requires precise knowledge of coupling rates to the cavity, which can be hard to calibrate and may depend on the Rabi frequency (similar to ΔI). In addition, it is difficult to experimentally verify whether the cavity modes are tracked accurately, as this would require cavity-state tomography.

In Fig. 2(d) we show a further improvement in validation error of the standard Bayesian filter after incorporating the two main effects of the resonator memory in the state update equations (based on Appendix L). We label this method as "BF+analytics" accordingly. In addition to a reduced measurement backaction, which was already included in the BF + numerics method described above, this method also takes into account the rotation of the measurement eigenstates resulting in a much better filter for fast qubit dynamics. This simple correction assumes that the Rabi drive changes the measured observable appreciably over the timescale of the resonator decay, effectively averaging it along a Rabi oscillation segment. This assumption is valid only for faster drives that can outpace the total measurement dephasing rate. Slower drives (with $\Omega < 2\Gamma_d$) will have their dynamics dominated by the quantum Zeno effect instead, where the measurement rapidly pins the state to an eigenstate mixture and prevents coherent Rabi evolution from occurring [64]. Thus, our expectation is that the BF + analytics method should perform less well than the standard steady-state filter for smaller drives within this quantum Zeno regime, which is supported by the data for the two drive strengths, $\Omega/2\pi \approx 0$ and 0.2 MHz, smaller than the quantum Zeno threshold of $2\Gamma_d/2\pi \approx 0.4$ MHz.

APPENDIX H: VERIFICATION OF HETERODYNE MEASUREMENT BACKACTION

A phase-preserving measurement technique gives rise to both informational and phase backaction. In Fig. 11, we verify that the undriven trajectories reconstructed by the LSTM indeed show the signatures predicted by Eqs. (E1)–(E3).

We analyze the measurement backaction for dynamics in the $x = 0$, $y = 0$, and $z = 0$ planes separately and the plots show good agreement with the expected measurement backaction, since we observe phase backaction in the xy plane and informational backaction in the xz and yz

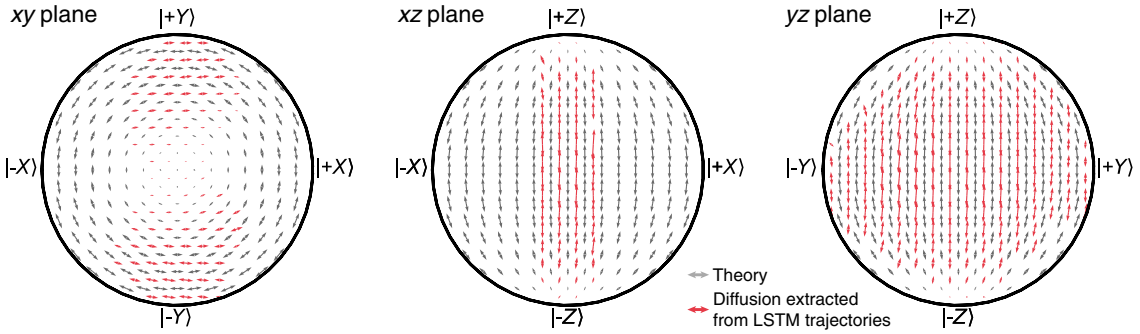

FIG. 11. Comparison between theoretical heterodyne measurement backaction (gray quiver maps) and reconstructed measurement backaction by the LSTM (red quiver maps) for an undriven qubit initially prepared in $|\pm Y\rangle$. The theory is obtained by setting z, y, $x = 0$ in the measurement backaction term of Eqs. (E1)–(E3) for backaction in the xy, xz, and yz plane, respectively. The reconstructed quiver maps do not cover the full xy and xz plane due to the limited measurement efficiency η but agree qualitatively with the theory.

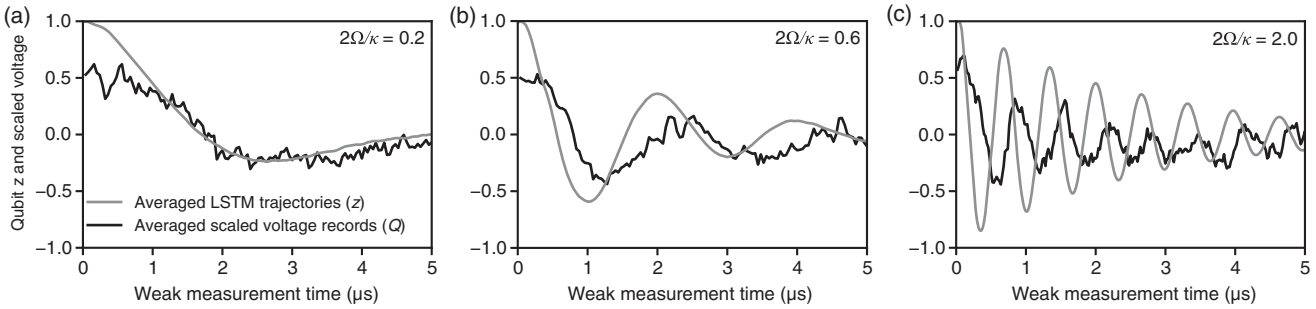

FIG. 12. Comparison between averaged voltage records and averaged z trajectories reconstructed by the LSTM for increasing Rabi frequencies: (a) $2\Omega/\kappa = 0.2$, (b) $2\Omega/\kappa = 0.6$, and (c) $2\Omega/\kappa = 2.0$. Both the trajectories and voltage records are averaged $M = 5 \times 10^3$ times.

planes. This suggests that the LSTM accurately predicts heterodyne measurement backaction from experimental observations, and that the parametrization of the measurement backaction correctly describes the experiment for $2\Omega/\kappa \ll 1$.

APPENDIX I: COMPARISON BETWEEN VOLTAGE RECORDS AND RECONSTRUCTED TRAJECTORIES

To further understand the breakdown of the Bayesian filter reconstruction demonstrated in Fig. 2(c) of the main text it is helpful to compare the averaged voltage records with the averaged qubit trajectories reconstructed by the LSTM. For fast qubit dynamics [Fig. 12(c)] we observe a large phase difference between the oscillations in the voltage record and those in $z(t)$, because the cavity memory delays photons escaping to the transmission line while the qubit rotates quickly. Conventional stochastic master equations or Bayesian filters that do not include the cavity memory assume photons measure the qubit state $z(t)$ instantaneously and, therefore, this phase difference signals the breakdown of those methods. Note that this phase difference already becomes noticeable at intermediate Rabi frequencies [Fig. 12(b)], consistent with the increase in rms error at $2\Omega/\kappa = 0.6$ shown in Fig. 2(d) of the main text.

APPENDIX J: EFFICIENCY CALIBRATION

To verify whether the LSTM correctly captures the measurement backaction we compare the extracted measurement rate from the LSTM with an independent calibration measurement, which extracts the separation $S = (\Delta V/\sigma)^2$ of conditional qubit histograms [Fig. 13(a)] after weak measurement of variable duration T_m. The characteristic measurement time measures how quickly we gain information about the qubit state in the presence of experimental inefficiencies, and is defined by the time it takes to reach a histogram separation of $\Delta V = 2\sigma$:

$$\tau_m = 4\left(\frac{dS}{dT_m}\right)^{-1}. \tag{J1}$$

In the undriven case where the coherent states in the cavity have reached steady state, the single quadrature measurement rate Γ_m is given by [61]

$$\Gamma_m = \frac{1}{2\eta\tau_m}, \tag{J2}$$

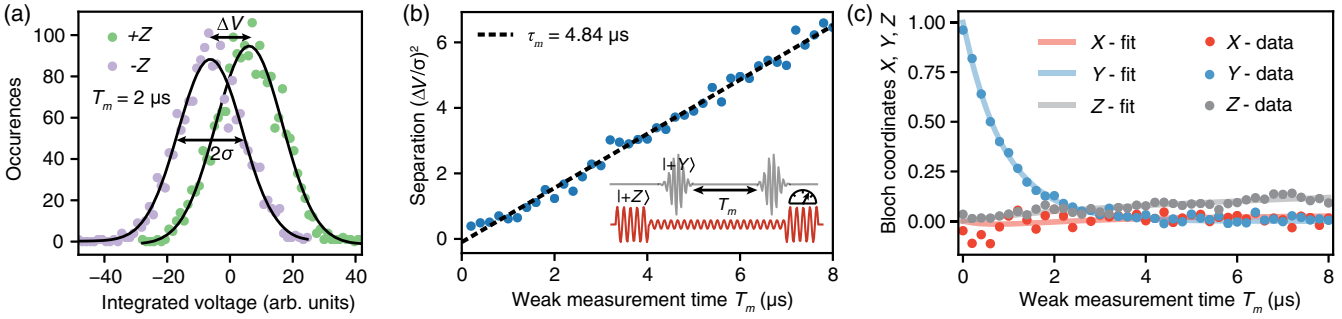

FIG. 13. (a) Histograms of the integrated voltage records conditioned on the projective measurement result after preparing the qubit in $|+Y\rangle$ and $T_m = 2.0$ μs of weak measurement. The histograms separate by an amount ΔV due to collapse of the wave function as a result of the weak measurement. (b) The normalized separation $S = (\Delta V/\sigma)^2$ increases linearly with the weak measurement time and the slope gives the characteristic measurement time τ_m. The inset shows a pulse sequence used for this calibration measurement. (c) To extract the efficiency we fit the exponential decay of the projective measurements in the Y basis (blue), which gives $T_d = \Gamma_d^{-1} = 0.91$ μs, or $\Gamma_d/2\pi = 0.175$ MHz.

where η is the total efficiency of the measurement chain. In addition, from the definition of the measurement rate we have $\eta = 2\Gamma_m/\Gamma_d$ [32], where the factor of 2 accounts for both quadratures and Γ_d is the decay rate of the ensemble average [Fig. 13(c)]. The results of the calibration are shown in Figs. 13(b) and 13(c), and from Eq. (J1) we find $\tau_m = 4.85 \pm 0.08\ \mu$s and $\Gamma_d/2\pi = 0.175$ MHz, which gives $\eta = 0.188 \pm 0.003$. This value agrees with the value found from analyzing the stochastic map at $\Omega = 0$ ($\eta = 0.185 \pm 0.002$). Note that this efficiency calibration is not valid when the coherent states of the cavity are not in equilibrium, for example, when the qubit is subject to a Rabi drive or when $\Gamma_d > \kappa$. In those cases the methods presented in Ref. [65] could offer a way to calibrate the efficiency, but this is beyond the scope of this work.

APPENDIX K: DERIVATION OF THE CORRECTIONS TO THE MEASUREMENT BACKACTION

1. Problem setup

A superconducting transmon is measured by coupling it to a strongly detuned readout resonator, which shifts and splits the resonator frequency by a contrast of 2χ between the two qubit-state-conditioned resonances. The ratio $2\chi/\kappa$ of this contrast to the resonance half-width $\kappa/2$ characterizes the distinguishability between resonances, and thus the amount of qubit information per resonator photon that can be stored in correlations. More precisely, given an input drive $\hat{d}(t)e^{-i\omega_d t}$ tuned to the midpoint between qubit resonances, photons in the resonator mode $\hat{a}$ encode this qubit information in their relative phase, which follows from the Heisenberg evolution of the resonator mode in the rotating frame of the drive, $\hat{a}'(t) = -(\kappa/2)[1 + i(2\chi/\kappa)\hat{z}(t)]\hat{a}(t) + \sqrt{\kappa}\hat{d}(t)$.

2. General derivation

When the qubit $\hat{z}$ is stationary and the drive $\langle\hat{d}\rangle = -i\varepsilon/\sqrt{\kappa}$ fluctuates around a constant mean, the resulting steady state $\hat{a}^{\mathrm{SS}} = -i\sqrt{\bar{n}}\exp(-i\hat{\phi})$ has a Lorentzian mean photon number $\bar{n} = |2\varepsilon/\kappa|^2/[1 + (2\chi/\kappa)^2]$ and qubit-dependent phase $\hat{\phi} = \arctan[(2\chi/\kappa)\hat{z}]$, with a maximum phase contrast of $\Delta\phi_{\max} = 2\arctan(2\chi/\kappa)$. A homodyne measurement aligned with the quadrature of maximum separation thus has a maximum amplitude contrast $\Delta\bar{a}_{\max} = 2\sqrt{\bar{n}}\sin(\Delta\phi_{\max}/2) = \sqrt{\bar{n}}(4\chi/\kappa)/\sqrt{1 + (2\chi/\kappa)^2}$ that sets the rate $\gamma_m = (\eta\kappa/2)|\Delta\bar{a}_{\max}|^2 = \eta(8\chi^2\bar{n}/\kappa)/[1 + (2\chi/\kappa)^2]$ at which maximally separated steady states can be distinguished by the photon amplitudes escaping the resonator at rate $\kappa/2$ and being successfully collected with efficiency η. The uncollected photons and residual qubit-resonator entanglement further contribute to a total qubit ensemble-dephasing rate due to measurement $\Gamma_m \approx \gamma_m/\eta$.

In the presence of a qubit drive, the resonator response additionally filters the evolution of $\hat{z}(t)$ to produce an effectively adiabatic response to a time-ordered geometric series of its delay averages. That is, with a similarly constant drive $\langle\hat{d}\rangle = -i\varepsilon/\sqrt{\kappa}$ and $t \gg 2/\kappa$ to let transients decay, the resonator evolution has the recurrence relation,

$$\hat{a}(t) = -i\frac{2\varepsilon}{\kappa} - i\frac{2\chi}{\kappa}\int_0^t \hat{z}(t-\tau)\hat{a}(t-\tau)\frac{\kappa e^{-\kappa\tau/2}d\tau}{2}, \tag{K1}$$

which implies

$$\hat{a}(t) = -i\frac{2\varepsilon}{\kappa}\mathcal{T}\sum_{n=0}^{\infty}\left[-i\frac{2\chi}{\kappa}\int_0^t \hat{z}(t-\tau)dP(\tau)\right]^n. \tag{K2}$$

The convolution kernel in the Green's function is an exponential probability distribution $dP(\tau) = \kappa e^{-\kappa\tau/2}d\tau/2$ over delay times τ, normalized as $\int_0^\infty dP(\tau) = 1$ with mean and variance both equal to the time constant $\tau_c = 2/\kappa$.

When $\hat{z}$ varies slowly on the timescale τ_c it can be approximately pulled outside the integral of Eq. (K1) to yield the standard steady-state solution but with a time-dependent phase $\hat{\phi}(t) = \arctan[(2\chi/\kappa)\hat{z}(t)]$ that adiabatically tracks the qubit evolution. The next-order approximation treats the evolution as approximately linear within the exponential envelope $\hat{z}(t-\tau) \approx \hat{z}(t) - \hat{z}'(t)\tau$, which additionally *delays* the response to the qubit by the mean delay time $\tau_c = 2/\kappa$ to produce the effective phase $\hat{\phi}(t) = \arctan[(2\chi/\kappa)\hat{z}(t-\tau_c)]$. For more rapid evolution, part of the evolution is averaged, thus reducing the measurement contrast while rotating the measurement basis.

3. Applying the theory to constant Rabi drive

In the case of a constant Rabi drive, the delay average in Eq. (K1) can be computed directly. Assuming dominant harmonic evolution $\dot{\hat{z}} = \Omega\hat{y}$, $\dot{\hat{y}} = -\Omega\hat{z}$, repeated integration by parts of the delay-averaged $\hat{z}(t)$ when $t \gg \tau_c$ yields a pair of geometric series defining an effective $\hat{z}_{\mathrm{eff}}(t)$ characterized by an adiabaticity parameter $(2\Omega/\kappa)$:

$$\begin{aligned}\hat{z}_{\mathrm{eff}}(t) &= \int_0^t \hat{z}(t-\tau)dP(\tau)\\ &= \frac{1}{1 + (2\Omega/\kappa)^2}\hat{z}(t) - \frac{(2\Omega/\kappa)}{1 + (2\Omega/\kappa)^2}\hat{y}(t)\\ &= \sqrt{\eta_{\mathrm{av}}}[\cos\theta\hat{z}(t) - \sin\theta\hat{y}(t)].\end{aligned} \tag{K3}$$

The averaging both attenuates the eigenvalue contrast of $\hat{z}$ by an efficiency factor η_{av} and rotates the observable coupled to the resonator by an angle θ in the yz plane. The tilt angle and efficiency are thus

$$\theta = \arctan(\Omega\tau_c) = \arctan(2\Omega/\kappa), \tag{K4}$$

$$\sqrt{\eta_{\mathrm{av}}} = \cos\theta = \frac{1}{\sqrt{1+(2\Omega/\kappa)^2}}. \tag{K5}$$

At longer times $t \gg \tau_c$ the geometric series in Eq. (K1) then yields the standard steady state, but with a phase angle that depends upon the effective delay-averaged observable $\hat{z}_{\mathrm{eff}}(t)$ that is rotated by θ and with eigenvalues reduced by $\hat{z}^2_{\mathrm{eff}} = \eta_{\mathrm{av}}$. This tilt can be understood equivalently as two simultaneous measurements along z and y with differing measurement rates $\gamma_z = \eta_{\mathrm{av}}\gamma_m \cos^2\theta$ and $\gamma_y = \eta_{\mathrm{av}}\gamma_m \sin^2\theta$, that compete to rotate the effective measurement poles.

In the main text we fit the tilt angle and measurement rate extracted from the LSTM to the expressions for θ and $\gamma_y + \gamma_z = \eta_{\mathrm{av}}\gamma_m$ from Eq. (K5). These results are shown in Fig. 4(c).

APPENDIX L: DERIVATION OF MEASUREMENT BACKACTION MAPS WITH CAVITY EFFECTS

To predict measurement backaction statistics shown in Fig. 4(b), it is convenient to analyze an unnormalized Bloch-vector description $\vec{s} = (y, z, p)$ restricted to the yz plane, where $p \in (0, 1]$ is the total probability and the standard Bloch coordinates are obtained after renormalization, $\vec{S} = \mathcal{N}(\vec{s}) = (y/p, z/p)$. Focusing on the informational part of the measurement, a collected signal scaled to a measured observable $\hat{z}$ will have the form $rdt = \langle\hat{z}\rangle dt + \sqrt{\tau}dW$ integrated over a small time bin dt, where $\tau = 1/2\eta\Gamma_m$ is a characteristic timescale, Γ_m is the induced ensemble-dephasing rate, $\eta \in [0, 1]$ is the efficiency, and dW with $\langle dW\rangle = 0$ and $\langle dW^2\rangle = dt$ is a Weiner increment of zero-mean Gaussian noise.

Observing such a signal r induces a partial state collapse that has the form of a hyperbolic boost matrix [66,67]:

$$\vec{s} \xrightarrow{(r,dt,\tau)} \hat{M}_z(r, dt, \tau)\vec{s}, \tag{L1}$$

$$\hat{M}_z(r, dt, \tau) \equiv \begin{bmatrix} 1 & 0 & 0 \\ 0 & \cosh(rdt/\tau) & \sinh(rdt/\tau) \\ 0 & \sinh(rdt/\tau) & \cosh(rdt/\tau) \end{bmatrix}. \tag{L2}$$

Similarly, the measurement of a rotated observable $\hat{z}(\theta)$ that is tilted by an angle θ in the yz plane and further attenuated by an efficiency η_2 is readily obtained using a rotation matrix:

$$\hat{M}_{z(\theta)}(r, dt, \tau, \eta_2) \equiv \hat{R}_x(\theta)\hat{M}_z(r, dt, \tau/\eta_2)\hat{R}_x^{-1}(\theta), \tag{L3}$$

$$\hat{R}_x(\theta) \equiv \begin{bmatrix} \cos\theta & \sin\theta & 0 \\ -\sin\theta & \cos\theta & 0 \\ 0 & 0 & 1 \end{bmatrix}. \tag{L4}$$

Computing the normalized state increment to linear order in rdt yields the expected stochastic state increment:

$$d\vec{S}(r, dt, \tau, \eta_2, \theta) \equiv \mathcal{N}(\hat{M}_{z(\theta)}\vec{s} - \vec{s}) = \begin{bmatrix} dy \\ dz \end{bmatrix}$$
$$= \eta_2 \frac{rdt}{\tau}\begin{bmatrix} -yz\cos\theta + (1-y^2)\sin\theta \\ -zy\sin\theta + (1-z^2)\cos\theta \end{bmatrix} + \mathcal{O}(rdt)^2. \tag{L5}$$

Ensemble averaging the signal r produces its mean $\langle r\rangle = \langle\hat{z}(\theta)\rangle = z\cos\theta + y\sin\theta$, yielding the mean backaction,

$$\langle d\vec{S}\rangle = \begin{bmatrix} \langle dy\rangle \\ \langle dz\rangle \end{bmatrix}$$
$$\approx \eta_2 \frac{dt}{\tau}(z\cos\theta + y\sin\theta)\begin{bmatrix} -yz\cos\theta + (1-y^2)\sin\theta \\ -zy\sin\theta + (1-z^2)\cos\theta \end{bmatrix}, \tag{L6}$$

from collapse, suppressing the distinct ensemble-dephasing terms appearing at second order, $\langle(rdt)^2\rangle = \tau dt + \mathcal{O}(dt^2)$.

To compute the variance of this informational backaction, it is convenient to decompose the random variable r into a convex mixture, $rdt = (+1dt + \sqrt{\tau}dW_+)p_+ + (-1dt + \sqrt{\tau}dW_-)(1 - p_+)$ of two independent zero-mean Gaussian random variables, dW_+ and dW_-, conditioned on definite eigenvalues $\pm$ of $\hat{z}(\theta)$. Because of the linear prefactor of r in Eq. (L5), the variances can be readily computed using the conditional variance rule: the variance of the mixture is the mixture of the conditional variances plus the variance of the mixture means, $\mathrm{Var}(Y) = \langle\mathrm{Var}(Y|X)\rangle_X + \langle(Y|X)^2\rangle_X - (\langle Y|X\rangle_X)^2$. Computing $\mathrm{Var}(rdt)$, keeping the lowest order in dt, and taking the square root then yields the standard deviation:

$$\mathrm{s.d.}(d\vec{S}) \equiv \sqrt{\mathrm{Var}(d\vec{S})}$$
$$\approx \sqrt{2\eta_2\frac{dt}{\tau}}\begin{bmatrix} -yz\cos\theta + (1-y^2)\sin\theta \\ -zy\sin\theta + (1-z^2)\cos\theta \end{bmatrix}. \tag{L7}$$

Note that the directionality of the standard deviation has been preserved when taking the square root.

Using the tilt angle $\theta = -\theta = -\arctan(2\Omega/\kappa)$ and the additional time-averaging efficiency factor $\eta_2 = \eta_{\mathrm{av}} = \cos^2\theta = 1/[1 + (2\Omega/\kappa)^2]$ outlined in the previous section produces the variance vector plots in Fig. 4(b) of the main text.

[1] C. W. Gardiner and M. J. Collett, *Input and Output in Damped Quantum Systems: Quantum Stochastic Differential Equations and the Master Equation*, Phys. Rev. A **31**, 3761 (1985).

[2] A. Barchielli, *Measurement Theory and Stochastic Differential Equations in Quantum Mechanics*, Phys. Rev. A **34**, 1642 (1986).

[3] L. Diósi, *Continuous Quantum Measurement and Itô Formalism*, Phys. Lett. A **129**, 419 (1988).

[4] V. P. Belavkin, *Quantum Continual Measurements and A Posteriori Collapse on CCR*, Commun. Math. Phys. **146**, 611 (1992).

[5] H. J. Carmichael, *Quantum Trajectory Theory for Cascaded Open Systems*, Phys. Rev. Lett. **70**, 2273 (1993).

[6] H. M. Wiseman and G. J. Milburn, *Quantum Theory of Field-Quadrature Measurements*, Phys. Rev. A **47**, 642 (1993).

[7] A. N. Korotkov, *Selective Quantum Evolution of a Qubit State due to Continuous Measurement*, Phys. Rev. B **63**, 115403 (2001).

[8] J. Gambetta, A. Blais, M. Boissonneault, A. A. Houck, D. I. Schuster, and S. M. Girvin, *Quantum Trajectory Approach to Circuit QED: Quantum Jumps and the Zeno Effect*, Phys. Rev. A **77**, 012112 (2008).

[9] A. N. Korotkov, *Quantum Bayesian Approach to Circuit QED Measurement*, *Quantum Machines: Measurement and Control of Engineered Quantum Systems* (Oxford University Press, Oxford, 2011), pp. 533–556.

[10] A. N. Korotkov, *Quantum Bayesian Approach to Circuit QED Measurement with Moderate Bandwidth*, Phys. Rev. A **94**, 042326 (2016).

[11] R. Vijay, D. H. Slichter, and I. Siddiqi, *Observation of Quantum Jumps in a Superconducting Artificial Atom*, Phys. Rev. Lett. **106**, 110502 (2011).

[12] M. Hatridge, S. Shankar, M. Mirrahimi, F. Schackert, K. Geerlings, T. Brecht, K. M. Sliwa, B. Abdo, L. Frunzio, S. M. Girvin, R. J. Schoelkopf, and M. H. Devoret, *Quantum Back-Action of an Individual Variable-Strength Measurement*, Science **339**, 178 (2013).

[13] L. Sun, A. Petrenko, Z. Leghtas, B. Vlastakis, G. Kirchmair, K. M. Sliwa, A. Narla, M. Hatridge, S. Shankar, J. Blumoff, L. Frunzio, M. Mirrahimi, M. H. Devoret, and R. J. Schoelkopf, *Tracking Photon Jumps with Repeated Quantum Non-Demolition Parity Measurements*, Nature (London) **511**, 444 (2014).

[14] U. Vool, S. Shankar, S. O. Mundhada, N. Ofek, A. Narla, K. Sliwa, E. Zalys-Geller, Y. Liu, L. Frunzio, R. J. Schoelkopf, S. M. Girvin, and M. H. Devoret, *Continuous Quantum Nondemolition Measurement of the Transverse Component of a Qubit*, Phys. Rev. Lett. **117**, 133601 (2016).

[15] K. W. Murch, S. J. Weber, C. Macklin, and I. Siddiqi, *Observing Single Quantum Trajectories of a Superconducting Quantum Bit*, Nature (London) **502**, 211 (2013).

[16] P. Campagne-Ibarcq, P. Six, L. Bretheau, A. Sarlette, M. Mirrahimi, P. Rouchon, and B. Huard, *Observing Quantum State Diffusion by Heterodyne Detection of Fluorescence*, Phys. Rev. X **6**, 011002 (2016).

[17] Q. Ficheux, S. Jezouin, Z. Leghtas, and B. Huard, *Dynamics of a Qubit While Simultaneously Monitoring Its Relaxation and Dephasing*, Nat. Commun. **9**, 1926 (2018).

[18] S. Hacohen-Gourgy, L. S. Martin, E. Flurin, V. V. Ramasesh, K. B. Whaley, and I. Siddiqi, *Quantum Dynamics of Simultaneously Measured Non-Commuting Observables*, Nature (London) **538**, 491 (2016).

[19] N. Roch, M. E. Schwartz, F. Motzoi, C. Macklin, R. Vijay, A. W. Eddins, A. N. Korotkov, K. B. Whaley, M. Sarovar, and I. Siddiqi, *Observation of Measurement-Induced Entanglement and Quantum Trajectories of Remote Superconducting Qubits*, Phys. Rev. Lett. **112**, 170501 (2014).

[20] A. Chantasri, M. E. Kimchi-Schwartz, N. Roch, I. Siddiqi, and A. N. Jordan, *Quantum Trajectories and Their Statistics for Remotely Entangled Quantum Bits*, Phys. Rev. X **6**, 041052 (2016).

[21] R. Vijay, C. Macklin, D. H. Slichter, S. J. Weber, K. W. Murch, R. Naik, A. N. Korotkov, and I. Siddiqi, *Stabilizing Rabi Oscillations in a Superconducting Qubit Using Quantum Feedback*, Nature (London) **490**, 77 (2012).

[22] G. de Lange, D. Riste, M. J. Tiggelman, C. Eichler, L. Tornberg, G. Johansson, A. Wallraff, R. N. Schouten, and L. DiCarlo, *Reversing Quantum Trajectories with Analog Feedback*, Phys. Rev. Lett. **112**, 080501 (2014).

[23] Z. K. Minev, S. O. Mundhada, S. Shankar, P. Reinhold, R. Gutiérrez-Jáuregui, R. J. Schoelkopf, M. Mirrahimi, H. J. Carmichael, and M. H. Devoret, *To Catch and Reverse a Quantum Jump Mid-Flight*, Nature (London) **570**, 200 (2019).

[24] L. S. Martin, W. P. Livingston, S. Hacohen-Gourgy, H. M. Wiseman, and I. Siddiqi, *Implementation of a Canonical Phase Measurement with Quantum Feedback*, Nat. Phys. **16**, 1046 (2020).

[25] S. Hacohen-Gourgy and L. S. Martin, *Continuous Measurements for Control of Superconducting Quantum Circuits*, Adv. Phys. X **5**, 1813626 (2020).

[26] R. Mohseninia, J. Yang, I. Siddiqi, A. N. Jordan, and J. Dressel, *Always-On Quantum Error Tracking with Continuous Parity Measurements*, Quantum **4**, 358 (2020).

[27] J. Atalaya, A. N. Korotkov, and K. B. Whaley, *Error-Correcting Bacon-Shor Code with Continuous Measurement of Noncommuting Operators*, Phys. Rev. A **102**, 022415 (2020).

[28] Y.-H. Chen and T. A. Brun, *Continuous Quantum Error Detection and Suppression with Pairwise Local Interactions*, Phys. Rev. Research **2**, 043093 (2020).

[29] W. P. Livingston, M. S. Blok, E. Flurin, J. Dressel, A. N. Jordan, and I. Siddiqi, *Experimental Demonstration of Continuous Quantum Error Correction*, Nat. Commun. **13**, 2307 (2022).

[30] K. Jacobs and D. A. Steck, *A Straightforward Introduction to Continuous Quantum Measurement*, Contemp. Phys. **47**, 279 (2006).

[31] A. Chantasri, J. Dressel, and A. N. Jordan, *Action Principle for Continuous Quantum Measurement*, Phys. Rev. A **88**, 042110 (2013).

[32] E. Flurin, L. S. Martin, S. Hacohen-Gourgy, and I. Siddiqi, *Using a Recurrent Neural Network to Reconstruct Quantum Dynamics of a Superconducting Qubit from Physical Observations*, Phys. Rev. X **10**, 011006 (2020).

[33] A. Graves, *Supervised Sequence Labelling with Recurrent Neural Networks* (Springer-Verlag, Berlin, 2012).

[34] G. Carleo, I. Cirac, K. Cranmer, L. Daudet, M. Schuld, N. Tishby, L. Vogt-Maranto, and L. Zdeborová, *Machine Learning and the Physical Sciences*, Rev. Mod. Phys. **91**, 045002 (2019).

[35] M. Hibat-Allah, M. Ganahl, L. E. Hayward, R. G. Melko, and J. Carrasquilla, *Recurrent Neural Network Wave Functions*, Phys. Rev. Research **2**, 023358 (2020).

[36] J. Carrasquilla, G. Torlai, R. G. Melko, and L. Aolita, *Reconstructing Quantum States with Generative Models*, Nat. Mach. Intell. **1**, 155 (2019).

[37] D. F. Wise, J. J. L. Morton, and S. Dhomkar, *Using Deep Learning to Understand and Mitigate the Qubit Noise Environment*, PRX Quantum **2**, 010316 (2021).

[38] M. Y. Niu, V. Smelyanskyi, P. Klimov, S. Boixo, R. Barends, J. Kelly, Y. Chen, K. Arya, B. Burkett, D. Bacon *et al.*, *Learning Non-Markovian Quantum Noise from Moiré-Enhanced Swap Spectroscopy with Deep Evolutionary Algorithm*, arXiv:1912.04368.

[39] J. Steinmetz, D. Das, I. Siddiqi, and A. N. Jordan, *Continuous Measurement of a Qudit Using Dispersively Coupled Radiation*, Phys. Rev. A **105**, 052229 (2022).

[40] E. Genois, J. A. Gross, A. Di Paolo, N. J. Stevenson, G. Koolstra, A. Hashim, I. Siddiqi, and A. Blais, *Quantum-Tailored Machine-Learning Characterization of a Superconducting Qubit*, PRX Quantum **2**, 040355 (2021).

[41] P. Rouchon and J. F. Ralph, *Efficient Quantum Filtering for Quantum Feedback Control*, Phys. Rev. A **91**, 012118 (2015).

[42] H. M. Wiseman and G. J. Milburn, *Open Quantum Systems, Quantum Measurement and Control* (publisher Cambridge University Press, Cambridge, England, 2009), p. 97147.

[43] D. Szombati, A. Gomez Frieiro, C. Müller, T. Jones, M. Jerger, and A. Fedorov, *Quantum Rifling: Protecting a Qubit from Measurement Back Action*, Phys. Rev. Lett. **124**, 070401 (2020).

[44] T. M. Stace and S. D. Barrett, *Continuous Quantum Measurement: Inelastic Tunneling and Lack of Current Oscillations*, Phys. Rev. Lett. **92**, 136802 (2004).

[45] F. Yan, S. Gustavsson, J. Bylander, X. Jin, F. Yoshihara, D. G. Cory, Y. Nakamura, T. P. Orlando, and W. D. Oliver, *Rotating-Frame Relaxation as a Noise Spectrum Analyser of a Superconducting Qubit Undergoing Driven Evolution*, Nat. Commun. **4**, 2337 (2013).

[46] J. Gambetta, A. Blais, D. I. Schuster, A. Wallraff, L. Frunzio, J. Majer, M. H. Devoret, S. M. Girvin, and R. J. Schoelkopf, *Qubit-Photon Interactions in a Cavity: Measurement-Induced Dephasing and Number Splitting*, Phys. Rev. A **74**, 042318 (2006).

[47] M. Naghiloo, A. N. Jordan, and K. W. Murch, *Achieving Optimal Quantum Acceleration of Frequency Estimation Using Adaptive Coherent Control*, Phys. Rev. Lett. **119**, 180801 (2017).

[48] A. Shabani, J. Roden, and K. B. Whaley, *Continuous Measurement of a Non-Markovian Open Quantum System*, Phys. Rev. Lett. **112**, 113601 (2014).

[49] S. Krastanov, K. Head-Marsden, S. Zhou, S. T. Flammia, L. Jiang, and P. Narang, *Unboxing Quantum Black Box Models: Learning Non-Markovian Dynamics*, arXiv:2009.03902.

[50] L. Banchi, E. Grant, A. Rocchetto, and S. Severini, *Modelling Non-Markovian Quantum Processes with Recurrent Neural Networks*, New J. Phys. **20**, 123030 (2018).

[51] T. Beucler, M. Pritchard, S. Rasp, J. Ott, P. Baldi, and P. Gentine, *Enforcing Analytic Constraints in Neural Networks Emulating Physical Systems*, Phys. Rev. Lett. **126**, 098302 (2021).

[52] A. van den Oord, S. Dieleman, H. Zen, K. Simonyan, O. Vinyals, A. Graves, N. Kalchbrenner, A. Senior, and K. Kavukcuoglu, *WaveNet: A Generative Model for Raw Audio*, arXiv:1609.03499.

[53] https://github.com/qnl/trajectories_lstm.

[54] J. Johansson, P. Nation, and F. Nori, *QuTiP2: A PYTHON Framework for the Dynamics of Open Quantum Systems*, Comput. Phys. Commun. **184**, 1234 (2013).

[55] D. I. Schuster, A. Wallraff, A. Blais, L. Frunzio, R.-S. Huang, J. Majer, S. M. Girvin, and R. J. Schoelkopf, *ac Stark Shift and Dephasing of a Superconducting Qubit Strongly Coupled to a Cavity Field*, Phys. Rev. Lett. **94**, 123602 (2005).

[56] S. Hochreiter and J. Schmidhuber, *Long Short-Term Memory*, Neural Comput. **9**, 1735 (1997).

[57] F. A. Gers, J. Schmidhuber, and F. Cummins, *Learning to Forget: Continual Prediction with LSTM*, Neural Comput. **12**, 2451 (2000).

[58] D. P. Kingma and J. Ba, *Adam: A Method for Stochastic Optimization*, arXiv:1412.6980.

[59] L. N. Smith, *Cyclical Learning Rates for Training Neural Networks*, in *Proceedings of the 2017 IEEE Winter Conference on Applications of Computer Vision (WACV)* (IEEE, Santa Rosa, CA, 2017), pp. 464–472, https://www.semanticscholar.org/paper/Cyclical-Learning-Rates-for-Training-Neural-Smith/37b5dfe87d82ba8f310155165d5bf841dc92dea2.

[60] N. Srivastava, G. Hinton, A. Krizhevsky, I. Sutskever, and R. Salakhutdinov, *Dropout: A Simple Way to Prevent Neural Networks from Overfitting*, J. Machine Learn. Res. **15**, 1929 (2014), https://jmlr.org/papers/v15/srivastava14a.html.

[61] J. Atalaya, S. Hacohen-Gourgy, I. Siddiqi, and A. N. Korotkov, *Correlators Exceeding One in Continuous Measurements of Superconducting Qubits*, Phys. Rev. Lett. **122**, 223603 (2019).

[62] S. J. Weber, A. Chantasri, J. Dressel, A. N. Jordan, K. W. Murch, and I. Siddiqi, *Mapping the Optimal Route between Two Quantum States*, Nature (London) **511**, 570 (2014).

[63] S. J. Weber, K. W. Murch, M. E. Kimchi-Schwartz, N. Roch, and I. Siddiqi, *Trajectoires Quantiques de Qubits Supraconducteurs*, C.R. Phys. **17**, 766 (2016).

[64] S. Hacohen-Gourgy, L. P. García-Pintos, L. S. Martin, J. Dressel, and I. Siddiqi, *Incoherent Qubit Control Using the Quantum Zeno Effect*, Phys. Rev. Lett. **120**, 020505 (2018).

[65] C. C. Bultink, B. Tarasinski, N. Haandbaek, S. Poletto, N. Haider, D. J. Michalak, A. Bruno, and L. DiCarlo, *General Method for Extracting the Quantum Efficiency of Dispersive Qubit Readout in Circuit QED*, Appl. Phys. Lett. **112**, 092601 (2018).

[66] A. N. Jordan and A. N. Korotkov, *Qubit Feedback and Control with Kicked Quantum Nondemolition Measurements: A Quantum Bayesian Analysis*, Phys. Rev. B **74**, 085307 (2006).

[67] T. L. Patti, A. Chantasri, L. P. García-Pintos, A. N. Jordan, and J. Dressel, *Linear Feedback Stabilization of a Dispersively Monitored Qubit*, Phys. Rev. A **96**, 022311 (2017).