

Using a Recurrent Neural Network to Reconstruct Quantum Dynamics of a Superconducting Qubit from Physical Observations

E. Flurin^{1,2,4,*}, L. S. Martin,^{1,2} S. Hacoheh-Gourgy,^{1,2,3} and I. Siddiqi^{1,2}

¹*Department of Physics, University of California, Berkeley, California 94720, USA*

²*Center for Quantum Coherent Science, University of California, Berkeley, California 94720, USA*

³*Department of Physics, Technion-Israel Institute of Technology, Haifa 32000, Israel*

⁴*Quantronics Group, SPEC, CEA, CNRS, Université Paris-Saclay, CEA-Saclay, 91191 Gif-sur-Yvette, France*

 (Received 14 February 2019; revised manuscript received 8 October 2019; published 9 January 2020)

At its core, quantum mechanics is a theory developed to describe fundamental observations in the spectroscopy of solids and gases. Despite these practical roots, however, quantum theory is infamous for being highly counterintuitive, largely due to its intrinsically probabilistic nature. Neural networks have recently emerged as a powerful tool that can extract nontrivial correlations in vast datasets. These networks routinely outperform state-of-the-art techniques in language translation, medical diagnosis, and image recognition. It remains to be seen if neural networks can be trained to predict stochastic quantum evolution without *a priori* specifying the rules of quantum theory. Here, we demonstrate that a recurrent neural network can be trained in real time to infer the individual quantum trajectories associated with the evolution of a superconducting qubit under unitary evolution, decoherence, and continuous measurement from physical observations only. The network extracts the system Hamiltonian, measurement operators, and physical parameters. It is also able to perform tomography of an unknown initial state without any prior calibration. This method has the potential to greatly simplify and enhance tasks in quantum systems such as noise characterization, parameter estimation, feedback, and optimization of quantum control.

DOI: [10.1103/PhysRevX.10.011006](https://doi.org/10.1103/PhysRevX.10.011006)

Subject Areas: Quantum Information

I. INTRODUCTION

Quantum mechanics breaks dramatically with classical intuition, contradicting determinism and introducing many highly counterintuitive concepts, such as contextuality, nonclassical correlations, and the uncertainty principle. Despite its abstract mathematical framework, quantum mechanics can be formulated operationally as an extended information theory [1], where the physical system is treated as a black box in which preparation and measurement combine to give the probabilities of experimental outcomes. The physical parameters are then estimated by averaging measurement outcomes on a large ensemble. The time evolution of the state of an isolated quantum mechanical system is governed by the Schrödinger equation. However, a realistic system cannot be isolated perfectly, and the coupling to an environment brings about qualitatively different behavior that cannot be accounted for via

the Schrödinger equation alone. If the system is monitored continuously, the dynamics of the system is perturbed by the inevitable backaction induced by measurement. Although the system's evolution under measurement is stochastic, the measurement record faithfully reports the perturbation of the system with respect to the unperturbed coherent evolution. Consequently, the observer's knowledge of the wave function can be updated using quantum filtering—the extraction of quantum information from a noisy signal. The stochastic time evolution of the wave function is the so-called quantum trajectory. Under certain approximations, this task can be performed by integrating the stochastic quantum master equation, provided that the Hamiltonian, dissipation, and measurement operators are precisely calibrated [2–5]. On the other hand, recurrent neural networks (RNNs) are a powerful class of machine-learning tools able to extract hidden correlations from large datasets [6]. They are most commonly applied to time-binned data and, as such, achieve excellent performance on difficult problems such as language translation [7] and speech recognition [8]. RNN training is driven by examples and performed without specifying dictionaries or linguistic rules. Interestingly, quantum filtering [9] can be seen as a similar task in which noisy experimental signals must be translated into meaningful quantum information. Last year,

*emmanuel.flurin@cea.fr

Published by the American Physical Society under the terms of the Creative Commons Attribution 4.0 International license. Further distribution of this work must maintain attribution to the author(s) and the published article's title, journal citation, and DOI.

various architectures of neural networks have been used in the realm of quantum physics for the prediction the theoretical quantum behavior of strongly correlated phases of matter [10–14], the design of efficient quantum error correction code [15], the decoding of large topological error correcting codes [16–18], and the optimization of dynamical decoupling schemes for quantum memories [19].

In this paper, we show that neural networks can be trained to predict stochastic quantum evolution from physical observations without specifying quantum mechanics *a priori*. We demonstrate that the RNN reproduces the stochastic quantum evolution for a continuously monitored superconducting qubit under a Rabi Hamiltonian. Rather than providing a black-box model, we use the neural network to robustly extract all physical parameters required for quantum filtering. Moreover, while RNNs are temporally oriented, they are routinely trained in both the forward and backward time ordering, so that the network may exploit both past and future information. In the present application, the use of past and future continuous measurement outcomes improves the estimation accuracy of quantum trajectories at a given time through a process called quantum smoothing [20,21]. We train a bidirectional RNN to perform a forward-backward analysis of trajectories, enabling quantum smoothing of predictions and the faithful tomography of an unknown initial state. By treating preparation and measurement on the same footing, the RNN structure highlights the time symmetry underlying the stochastic quantum evolution.

II. EXPERIMENTAL SYSTEM

Our experiment consists of a superconducting transmon qubit [22] dispersively coupled to a superconducting waveguide cavity [23]. In the interaction picture and rotating wave approximation, our system is described by the Hamiltonian $H = H_{\text{int}} + H_R$:

$$H_{\text{int}} = \frac{\hbar\chi}{2} a^\dagger a \sigma_Z, \quad (1)$$

$$H_R = \frac{\hbar\Omega_R}{2} \sigma_X, \quad (2)$$

where \hbar is the reduced Planck's constant, a^\dagger (a) is the creation (annihilation) operator for the cavity mode, and $\sigma_{X,Y,Z}$ are qubit Pauli operators where the qubit ground state is the eigenmode of the σ_Z operator associated for the eigenvalue -1 . H_R describes a microwave drive at the qubit transition frequency which induces unitary evolution of the qubit state characterized by the Rabi frequency $\Omega_R = 2\pi \times 0.80$ MHz. H_{int} is the interaction term, characterized by the dispersive coupling rate $\chi = -2\pi \times 0.18$ MHz. This term describes a qubit state-dependent frequency shift of the cavity, which we use to perform a quantum state measurement of our qubit.

The cavity is coupled to the transmission line at a rate $\kappa = 2\pi \times 7.2$ MHz. A microwave tone that probes the cavity near its resonance frequency will acquire a qubit state-dependent phase shift. If the measurement tone is very weak, quantum fluctuations of the electromagnetic mode fundamentally obscure this phase shift, resulting in a partial or weak measurement of the qubit state [2]. The strength of the measurement is set such that the qubit-induced dephasing rate γ_ϕ is on par with the Rabi timescale, here $\gamma_\phi = 1.1 \mu\text{s}^{-1}$. We use a near-quantum-limited parametric amplifier [24] to amplify the quadrature of the reflected signal which is proportional to the qubit state-dependent phase shift. After further amplification, we digitize the signal in 40 ns time steps, yielding a measurement record V_t .

We begin each run of the experiment by heralding the ground state of the qubit using the above readout technique. We then prepare the qubit along one of the six cardinal points of the Bloch sphere by applying a preparation pulse. Next, a measurement tone at the cavity frequency of 6.666 GHz continuously probes the cavity for a variable time T between 0 and 4 μs , which weakly measures the qubit in the σ_Z basis. Concurrently, we apply the Rabi Hamiltonian H_R . Finally, we apply pulses to perform qubit rotations and a projective measurement, yielding a single shot measurement of the desired qubit operator σ_X , σ_Y , or σ_Z .

III. QUANTUM TRAJECTORIES

To allow the neural network to operate as generally as possible, we formulate system inputs and outputs symmetrically and avoid passing it objects such as a wave function that encodes information about the structure of quantum theory. The role of the wave function in quantum mechanics is to provide the probability of a measurement outcome y_t given the preparation and evolution of the system at earlier times $P(y_t|y_0)$. In the case of a continuously monitored quantum bit, the preparation and measurement outcome are each a binary variable $y_0, y_t \in \{0, 1\}$ extracted through a projective readout performed at the initial and final times, respectively; the preparation and measurement configurations, labeled a and b , encode microwave pulses performing qubit rotations for state preparation and tomography, respectively, in the X , Y , and Z basis. The stochastic measurement record $\{V_t\}$ is collected with a high quantum efficiency parametric amplifier during the qubit evolution.

Quantum trajectory theory describes how an observer's state of knowledge evolves given a measurement record [25]. Therefore, quantum trajectories are specified by $P(y_t|y_0, a, b, V_0 \dots V_t)$, the probability of measuring the outcome y_t with the measurement parameter b given the initial measurement y_0 in the preparation parameter a and the stochastic measurement outcome up to a time t . Tracking this quantum evolution can be understood as a translation of the measurement records into a quantum state evolution.

Figure 2(a) shows the distribution of measurement records obtained for the preparation setting ($y_0 = 0$, $a = Z$).

Quantum trajectories are typically extracted from a continuous measurement by integrating the stochastic master equation (SME) governing the evolution of the density matrix ρ_t :

$$d\rho_t = \underbrace{\left(i[H_R, \rho_t] + \mathcal{L} \left[\sqrt{\frac{\gamma_\phi}{2}} \sigma_Z \right] \rho_t \right)}_{\text{deterministic evolution}} dt + \underbrace{\sqrt{\eta} \mathcal{H} \left[\sqrt{\frac{\gamma_\phi}{2}} \sigma_Z \right] \rho_t}_{\text{backaction}} dw_t, \quad (3)$$

where \mathcal{L} is the Lindblad superoperator describing the qubit dephasing induced by the measurement of strength γ_ϕ , \mathcal{H} is a measurement superoperator describing the backaction of the measurement on the quantum state for a quantum efficiency η , and dw_t is a Gaussian-distributed variable with variance dt extracted from the measurement record normalized appropriately using

$$dw_t = \left(V_t - 2\sqrt{\eta} \text{Tr} \left[\rho_t \sqrt{\frac{\gamma_\phi}{2}} \sigma_Z \right] \right) dt. \quad (4)$$

The probability distribution for the projective outcome is then given by the Born rule $P_{X,Y,Z}(t) = P(y_t|y_0, a, b = X, Y, Z, V_0 \dots V_t) = (\text{Tr}[\rho_t \sigma_{X,Y,Z}] + 1)/2$. The integrated stochastic master equation provides faithful predictions when experimental parameters are precisely known from independent calibration under the assumption that the cavity decay rate is much larger than the qubit measurement rate $\kappa \gg \gamma_\phi$. Figure 2(a) shows two representative trajectories extracted from the measurement records based on the stochastic master equation.

IV. RECURRENT NEURAL NETWORK

Based solely on a large set of labeled examples $(y_t, y_0, a, b, \{V_\tau\})$ directly extracted from the experimental system, we now demonstrate that the network can be trained to predict the probability $P(y_t|y_0, a, b, V_0 \dots V_t)$ of observing the measurement outcome $y_t \in \{0, 1\}$ given the history of the quantum evolution accessible to the observer—in other words, the best knowledge of the qubit wave function.

We use a long short-term memory recurrent neural network [26] schematically depicted in Fig. 1(b). These typically consist of a layer of n virtual neuronlike nodes recurrently updated in time. The state of the neuron's layer at a time t is encoded in an n -dimensional vector \vec{h}_t . It is computed as a weighted linear combination of the neuron's layer state at a previous time $t-1$ combined with the measurement record at a time t and passed through a

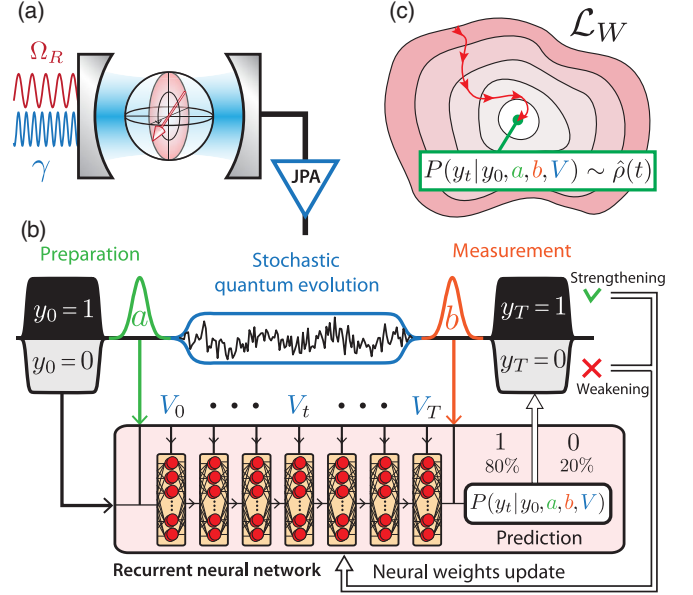


FIG. 1. Recurrent neural network training from physical observations. (a) Schematic of the superconducting qubit dispersively coupled to a microwave cavity monitored by a high quantum efficiency Josephson parametric amplifier (JPA). The qubit is simultaneously driven on resonance at a Rabi rate Ω_R and dispersively monitored with a strength γ_ϕ near the cavity resonance frequency. (b) Data collected from the experimental system, comprising the preparation, measurement outcomes, and continuous measurement record of the qubit, are directly streamed to a RNN, which provides a prediction of the measurement outcome. The weights of the RNN are updated at each iteration through stochastic gradient descent. (c) The stochastic gradient descent aims at minimizing the cross-entropy loss function \mathcal{L}_W , which evaluates the distance between the prediction and the measurement outcome.

nonlinear activation function ϕ such that $\vec{h}_t = \phi(W_h \cdot \vec{h}_{t-1} + V_t \vec{W}_V + \vec{B}_h)$, where W and B are the weights of the connections between the neurons and the biases, respectively, which are determined during the training stage. The probability $P_b(y_t)$ of getting the outcome y given the measurement setting b is computed at each time step as a linear combination of the neuron layer state passed through the activation function given by $P(y_t|y_0, a, b, V_0 \dots V_t) = \sigma(\vec{W}_b \cdot \vec{h}_t + \vec{B}_b)$. The preparation settings a and the initial qubit state (input bit y_0) are specified in the initial state of the neuron layer. The neural network is trained to minimize a loss function \mathcal{L} by strengthening or weakening connections between neuron layers encoded in the weight matrices $W_{h,v,b}$, as shown in Fig. 1(b). The cross-entropy loss function $\mathcal{L}^b = -y_T \log P(y_T|y_0, a, b, V_0 \dots V_T) - (1 - y_T) \log (1 - P(y_T|y_0, a, b, V_0 \dots V_T))$ is minimized when the prediction $P(y_t|y_0, a, b, V_0 \dots V_t)$ and the distribution of experimental outcomes y_T for a given measurement setting b match. Crucially, the function implemented by the neural network is differentiable, and, therefore, the weight

matrices can be updated at each iteration of the training by differentiating the loss function and applying a gradient-descent minimization step: $W \leftarrow W - \xi \langle \partial \mathcal{L}^b / \partial W \rangle$, where ξ is the learning rate. The training process ends once the weight matrices W converge toward a minimum of the loss function. The effectiveness of neural networks lies in their ability to converge toward a minimum of a very high-dimensional nonlinear loss landscape through gradient backpropagation as illustrated in Fig. 1(c).

V. TRAINING

The long short-term memory recurrent neural network comprises 64 neurons with rectified linear unit activation functions. This specific RNN architecture evades the exploding or vanishing gradient problem of standard RNN architectures, improving the learning of long-term dependencies [27]. The neural network is implemented with the Tensorflow library [28] developed by Google and optimized for a graphics processing unit (Nvidia Tesla K80 GPU), which enables a speedup of the training.

The data are fed to the network in batches, each containing 1024 measurement records on which a step of the gradient descent is performed using the so-called Adam optimization algorithm [29]. The measurement records are split into two datasets. 1.5×10^6 traces are used for the training, and 5×10^5 randomly chosen traces are used for evaluation and displayed in the manuscript. The training data can be reinjected several times to the network in order to improve the model accuracy. Each of these training cycles corresponds to a training epoch. In practice, up to ten training epochs have been performed. At each training epoch, the learning rate is lowered from 1×10^{-3} to 1×10^{-6} . In order to improve the training robustness, 30% of the neurons are dropped out randomly during the first epoch. The fraction of dropped out neurons is gradually lowered to 0 with each subsequent training epoch. This method prevents the network from overfitting and helps the generalization abilities of the model [30]. Note that the training quality does not strongly depend on the details of these parameters. A key feature of the training is that it can be performed in real time directly from data collected from the experimental system.

A common practice in the field of deep learning is to maximize the number of neurons and, hence, the number of free parameters, in order to saturate the computation power for a given training time [31]. In this work, the training time is given by the repetition rate of the experiment (0.5 ms per trace) which is limited by the qubit relaxation time ($T_1 = 60 \mu\text{s}$). Therefore, the number of neurons (64) is picked such that the training can be performed in 0.8 ms per trace on average. The 0.3 ms discrepancy between the experiment and training accounts for data processing, saving, and data transfer to the cloud. Therefore, the 2×10^6 traces are produced and fed to the RNN in 20 min. Six

preparation settings ($y_0 \in \{0, 1\}$, $a \in \{X, Y, Z\}$) and six measurement settings ($y_T \in \{0, 1\}$, $b \in \{X, Y, Z\}$) are used. In practice, we perform the preparation and measurement with the following rotations of the qubit— $R_{\pi/2}^Y$, $R_{-\pi/2}^Y$, $R_{\pi/2}^X$, $R_{-\pi/2}^X$, R_0^X , and R_π^X —which correspond to the cardinal points of the Bloch sphere. The associated preparation labels (y_0, a) and measurement labels (y_T, b) are then given, respectively, by (y, X) , (\bar{y}, X) , (y, Y) , (\bar{y}, Y) , (y, Z) , and (\bar{y}, Z) with $\bar{y} = 1 - y$. The total time evolution is varied over 20 values within $4 \mu\text{s}$ ($T \in [0, 4]$), and the measurement record $\{V_t\}$ is acquired during the qubit evolution with a sampling time of 40 ns. Once the training phase is completed, the RNN returns the prediction $P(y_t | y_0, a, b, V_0 \dots V_t)$, which can be interpreted as the probability of measuring the qubit at a time t along the measurement axis $b = X, Y$, and Z .

VI. VALIDATION

Once the RNN is trained, the predictions of the measurement outcomes form an ensemble of trajectories for each of the measurement settings as shown in Fig. 2(b). The predictions of the neural network are in good agreement with the representative trajectories integrated from the SME. In this section, we show that the accuracy of the quantum trajectory predicted by the SME and the RNN can be precisely quantified, and we demonstrate that the RNN predictions are more accurate than the ones inferred from the SME with respect to the projective measurement outcomes. The quality of the training can be evaluated self-consistently on the evaluation dataset (not used during the training). This method has been previously used to benchmark the prediction of the stochastic master equation [2–5]. We select the subset of the trajectories leading to the same prediction p within a small δ such that $\mathcal{S}_p = \{y_T \text{ such that } P(y_T | y_0, a, b, V_0 \dots V_T) \in [p - \delta, p + \delta]\}$. Figures 2(c) and 2(d) display the agreement between the ensemble of trajectories ending in $p \pm \delta = 0.85 \pm 0.015$ and the histogram of the final measurement value. If the prediction is accurate, it should agree with the final tomographic measurement average on the subset \mathcal{S}_p , defined as $\langle y \rangle_{\mathcal{S}_p} = \mathcal{N}_p^{-1} \sum_{y \in \mathcal{S}_p} y$ with \mathcal{N}_p the number of trajectories in \mathcal{S}_p , such that $\langle y \rangle_{\mathcal{S}_p} = p$. The overall agreement between the prediction and the tomography values can be quantified as an average error $\epsilon = \sum_p (\mathcal{N}_p / \mathcal{N}) \times (\langle y \rangle_{\mathcal{S}_p} - p)^2$, where \mathcal{N} is the total number of trajectories. The RNN prediction gives an average error lower than 10^{-2} for the all-measurement axis. As a comparison, using the same evaluation dataset, the prediction of the stochastic master equation (numerically propagated using the methods of Ref. [32]) based on the independently calibrated experimental parameters gives a higher average error along the Y and Z axes. Such a discrepancy can be attributed to small calibration errors and experimental drifts. This self-

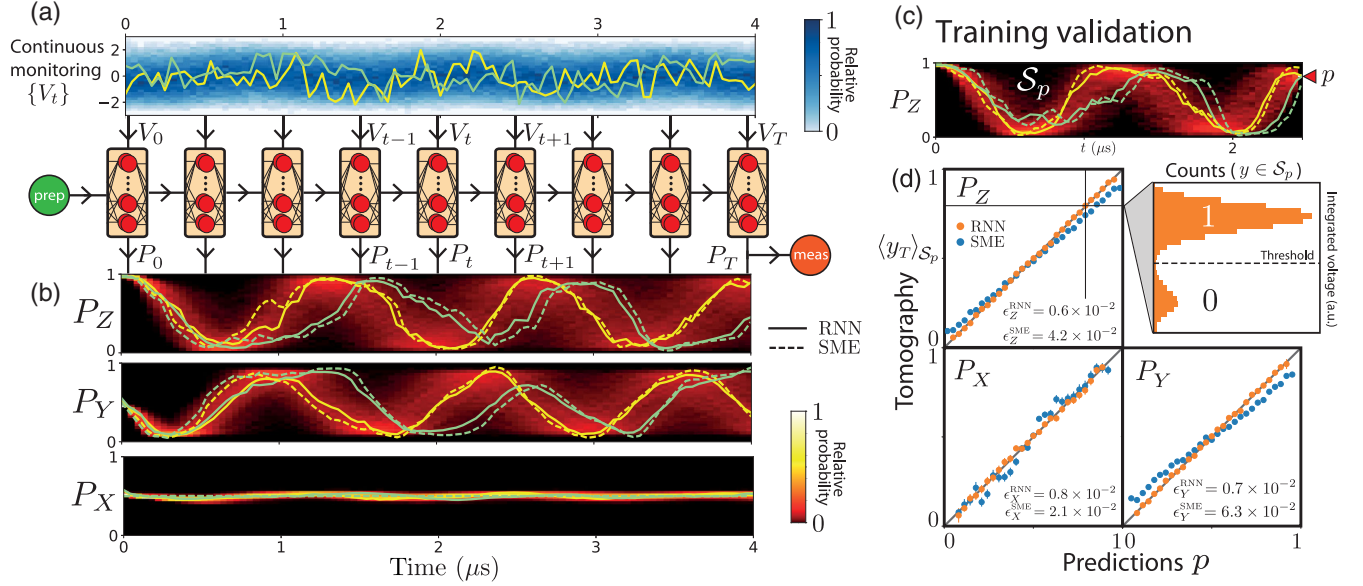


FIG. 2. RNN prediction of the quantum evolution. (a) Blue-scale histogram of the normalized measurement records extracted from the experiment; traces plotted in color show two representative instances. (b) Red-scale histograms of RNN prediction for the measurement basis $b = X, Y$, and Z in the driven case, beginning from $y_0 = 1$ in the preparation basis $a = X$. Traces plotted in color show the two representative predictions from the measurement records by the RNN (plain line) and by the stochastic master equation (dashed line). (c) Training validation; red-scale histogram of the RNN prediction leading to $p \pm \delta = 0.85 \pm 0.015$ at $T = 2.5 \mu\text{s}$ indicated by the red maker. This subset of trajectories belongs to the validation subset \mathcal{S}_p . (d) Comparison of the RNN prediction with the tomography. The tomography $\langle y_T \rangle_{\mathcal{S}_p}$ (y axis) is given by the projective measurement outcome along X, Y , and Z quadratures averaged over the subset of trajectories \mathcal{S}_p leading the prediction p (x axis). The orange dots correspond to the RNN predictions, and the blue dots correspond to the SME predictions. The average errors $\epsilon_{X,Y,Z}$ shown in the plots are used to quantify the agreements between the prediction and the tomography; it is calculated based on the method explained in the validation section of the main text. Inset: Histogram corresponding to the integrated voltage during the projective measurement for the quantum trajectory subset \mathcal{S}_p for $p = 0.85 \pm 0.015$. The threshold is used to discriminate the 0 and 1 states, leading to a tomography prediction $\langle y_T \rangle_{\mathcal{S}_p} = 0.85$ in perfect agreement with the RNN prediction. The SME prediction indicates \mathcal{S}_p .

consistent evaluation demonstrates the prediction power of the trained RNN and its robustness against calibration errors of physical parameters.

VII. BIDIRECTIONAL RNN

RNNs are inherently time oriented; the prediction $P(y_t|y_0, a, b, V_0 \dots V_t)$, at a time t , depends only on the measurement record at earlier times. A common feature used to improve the prediction power of a RNN, for a translation application, in particular, is to combine the prediction of two RNNs trained, respectively, forward and backward in time, exploiting the same data in both directions [6]. The forward prediction provides the trajectory given the past measurement record ($V_0 \rightarrow V_t$) and the preparation settings (y_0, a): $P^{\Rightarrow}(y_t) = P(y_t|y_0, a, b, V_0 \dots V_t)$, while the backward prediction provides the trajectory given the “future” measurement record ($V_T \rightarrow V_t$) played backward and the measurement settings (y_T, b): $P^{\Leftarrow}(y_t) = P(y_t|y_T, a, b, V_T \dots V_t)$. As shown in Fig. 3(a), the RNN provides an ensemble of backward trajectories. The accuracy of backward predictions is evaluated using the same validation method as the forward prediction;

the subset of backward trajectory \mathcal{S}_p giving the same prediction p must agree on average with the preparation measurement such that $\langle y_0 \rangle_{\mathcal{S}_p} = p$. The accuracy of the backward prediction is shown in Fig. 3(b), where the average errors for the preparation settings X, Y , and Z for the backward predictions are $\epsilon_X^{\Leftarrow} = 1.1 \times 10^{-2}$, $\epsilon_Y^{\Leftarrow} = 0.9 \times 10^{-2}$, and $\epsilon_Z^{\Leftarrow} = 0.7 \times 10^{-2}$, respectively, and the overall accuracy is comparable to the forward prediction. Remarkably, the backward and forward predictions do not necessarily agree at a given t ; indeed, these predictions are based on distinct parts of the measurement records. They provide complementary information from the past and future evolution of the system. These predictions can, therefore, be combined to enhance the knowledge of the quantum state based on the full measurement record. Backward-forward analysis is a well-established postprocessing method with recurrent neural network [6] as well as hidden Markov chain methods [33]. Time-reversal symmetry underlies quantum evolution and exchanges the role of state preparation and state measurement [34]. In a sense, backward-forward analysis naturally translates into the quantum regime as the prediction and retrodiction of




FIG. 3. RNN prediction and retrodiction of the quantum evolution. (a) Red-scale histograms of the RNN prediction for the measurement basis $b = Y$ and Z in the driven case beginning from $y_0 = 1$ in the preparation basis $a = Z$. Traces plotted in color show the two representative predictions from the measurement records by the RNN (plain line) and by the stochastic master equation (dashed line). (b) Blue-scale histogram of the normalized measurement records extracted from the experiment; traces plotted in color show representative instances. (c) Red-scale histograms of the RNN retrodiction for the same measurement record. Traces plotted in color show the two representative retrodictions from the measurement records by the RNN. (d) Retrodiction validation: Comparison of the RNN backward prediction with the tomography. The tomography $\langle y_0 \rangle_S$ (y axis) is given by the initial projective measurement outcome along X , Y , and Z quadratures averaged over the subset of trajectories S_p leading to the prediction p (x axis). The blue dots correspond to the RNN predictions. The average errors $\epsilon_{X,Y,Z}^{\leftarrow}$ shown in the plots are used to quantify the agreements between the prediction and the tomography; it is calculated based on the method explained in the validation section of the main text. (e) Red-scale histograms of smoothed RNN predictions based on the forward-backward analysis given by Eq. (5) for the same measurement records.

quantum trajectories [35–37]. Quantum predictions and retrodictions can be combined based on quantum smoothing techniques [20,21] enabling the enhancement of physical parameter estimation [38,39]. The forward and backward predictions can be combined into a smoothed prediction by

$$P^{\leftrightarrow}(y_t) = \frac{P^{\leftarrow}(y_t)P^{\rightarrow}(y_t)}{P^{\leftarrow}(y_t)P^{\rightarrow}(y_t) + [1 - P^{\leftarrow}(y_t)][1 - P^{\rightarrow}(y_t)]}. \quad (5)$$

As depicted in Fig. 3(c), the smoothed trajectories combine the backward and forward information such that it dismisses the least informative predictions [$P^{\leftarrow}(y_t), P^{\rightarrow}(y_t) \sim 0.5$] and strengthens the most informative ones [$P^{\leftarrow}(y_t), P^{\rightarrow}(y_t) \sim 0/1$]. By removing ambiguities in the qubit evolution, we access information which is blurred by statistical uncertainties in the standard approach, and we observe an improved temporal resolution on the qubit trajectory *a posteriori*. The forward-backward analysis

demonstrates how bidirectional RNNs naturally combine causal and anticausal correlations hidden in the measurement records.

VIII. INITIAL STATE ESTIMATION

The role of the preparation (y_0, a) and measurement (y_T, b) are treated symmetrically in the forward and backward prediction. Hence, while the forward RNN predicts the outcome of the final projective measurement, the backward RNN provides an estimation of the initial state of the system given the measurement record. These predictions can be, therefore, exploited to perform initial state tomography; this task is reminiscent of the enhanced readout discrimination by machine learning demonstrated in Ref. [40]. For the state estimation, we do not specify the final projective measurement, and we initialize the backward network with a maximally unknown state [$P^{\leftarrow}(y_T) = 0.5$ for X , Y , and Z]. Each backward trajectory provides up to one bit of information about the initial state [41]. Combining this information using maximum-likelihood

methods allows for reconstructing the initial state \vec{P}_0 . Here, the optimization consists in minimizing the following likelihood function over the probability of the initial state following Ref. [42]:

$$P_0(y_0|a) = \underset{P_0}{\operatorname{argmin}} \left(\sum_n |P_0 - P(y_0|a, V_T^{(n)} \dots V_0^{(n)})|^2 \right). \quad (6)$$

As shown in Fig. 4(a), we find an agreement between the initial state estimation and preparation within the 95% confidence interval estimated with bootstrapping method. It demonstrates that, despite the nontrivial dynamics resulting from the interplay between the Rabi evolution and the continuous monitoring, the combination of RNN backward predictions performs as a faithful qubit state tomography.

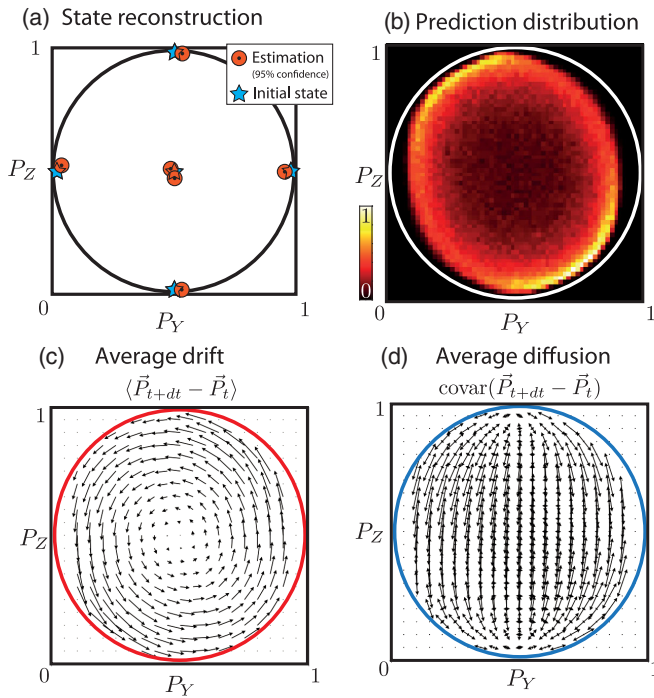


FIG. 4. Parameter estimation of the quantum master equation and initial state tomography. (a) State estimation. Estimation of six initial state preparations (red circles) using maximum-likelihood estimation on backward RNN predictions (approximately 20 000 trajectories each) initialized from an undetermined projective measurement outcome; the circle radius gives the 95% confidence interval extracted from bootstrapping methods. (b) Distribution of the RNN predictions in the Y and Z measurement basis for all time. (c) Average drift of individual trajectories in the Bloch sphere: the vector map of the averaged evolution of RNN predictions in the Y and Z measurement basis between two consecutive time steps. This map captures the Hamiltonian evolution and the Lindbladian dissipation. (d) Average diffusion of individual trajectories in the Bloch sphere: computed vector map associated with the covariance of the prediction between two consecutive time steps in the Y and Z measurement basis. This map captures the measurement-induced backaction.

IX. PARAMETER ESTIMATION

The trajectories predicted by the trained RNN can be exploited to estimate the physical parameters of the experimental system. In Fig. 4(b), we plot the distribution of the forward RNN prediction in the Y, Z plane for all times. This distribution exhibits a tilted ellipse shape within the Bloch sphere (white circle); the great axis of the ellipse is along the Z axis, showing that the quantum trajectories tend to collapse toward the poles of the Bloch sphere, corresponding to the pointer states of the measurement operator. In the equatorial plane, the distribution is squeezed, indicating that the quantum state experiences more dephasing and loses purity. These features are generated by the behavior of the system over individual time steps, which is captured by both the neural network predictions and the underlying physical model. The stochastic master equation has two main contributions [25]; on one hand, the deterministic evolution encodes the Hamiltonian evolution along with the decoherence, while, on the other hand, the measurement backaction describes the update of the quantum state given the stochastic measurement record. By performing a least-squares fit of a master equation model to the deterministic and diffusive components of the forward RNN prediction, we are able to reconstruct the physical parameters associated with the stochastic master equation describing the quantum evolution under continuous measurement.

The deterministic component of the evolution can be extracted from the forward prediction of the RNN by evaluating the average drift of individual trajectories. We compute the ensemble-averaged prediction change over intervals of 40 ns, $d\vec{P} = \langle \vec{P}_{t+1} - \vec{P}_t \rangle$ with $\vec{P}_t = [P_X(y_t), P_Y(y_t), P_Z(y_t)]$, versus the position on the Bloch sphere depicted in Fig. 4(c). We observe a vector map indicating an average drift of trajectories in the Bloch sphere. The drift corresponds to a rotation of the qubit state along the X axis of the Bloch sphere and an additional squeezing dynamics along the Z axis. The average dynamics is well reproduced by the deterministic evolution of the stochastic master equation in Eq. (3) for a Rabi Hamiltonian $H_R = \hbar\Omega_R(\sigma_x/2)$ with a Rabi frequency $\Omega_R/2\pi = 0.82 \pm 0.02$ MHz and a measurement-induced dephasing operator $\sqrt{\gamma_\phi/2}\sigma_Z$ with a rate $\gamma_\phi = 1.1 \pm 0.05 \mu\text{s}^{-1}$. The measurement-induced disturbance can also be extracted from the prediction of the RNN by evaluating the average diffusion of the individual trajectories [4]. We compute the covariance matrix associated with the prediction change over intervals of 40 ns, $dP^2 = \operatorname{covar}(\vec{P}_{t+1} - \vec{P}_t)$. The diffusion vector map is given by the eigenvectors of the covariance matrix weighted by its eigenvalues versus the position in the Bloch sphere as depicted in Fig. 4(b). This vector map describes the magnitude and the direction of the disturbance induced by the measurement in the Bloch sphere. We observe that the disturbance is maximal along

TABLE I. Summary of the parameter estimation based on RNN predictions.

	Calibration	RNN estimations
$\Omega_R/2\pi$	0.80 MHz	0.82 ± 0.02 MHz
γ_ϕ	$1.1 \mu\text{s}^{-1}$	$1.1 \pm 0.05 \mu\text{s}^{-1}$
γ_m	$0.40 \mu\text{s}^{-1}$	$0.40 \pm 0.01 \mu\text{s}^{-1}$

the equatorial plane of the Bloch sphere and vanishes at the poles. The diffusion is well reproduced by the deterministic evolution of the stochastic master equation in Eq. (3). From this map, we extract a measurement rate of $\gamma_m = 0.40 \pm 0.01 \mu\text{s}^{-1}$ along the Z axis of the Bloch sphere. The quantum efficiency of our measurement defined as the ratio of the measurement-induced dephasing and the measurement rate gives $\eta = \gamma_m/\gamma_\phi = 36\%$. Note that the quantum efficiency is usually challenging to estimate and requires several steps of calibrations. As summarized in Table I, the estimated experimental parameters differ slightly from the calibrations, which is attributed to residual detuning of the Rabi drive with respect to the qubit frequency.

X. CONCLUSION

We demonstrate that a recurrent neural network can be trained to provide a model-independent prediction of the outcome of fully general quantum evolution based only on a physical observation, at the level of a single qubit. The ensemble of predictions can be compared to quantum models such as the stochastic master equation to extract physical parameters without additional calibration. By considering causal and retrocausal evolution, we show that initial state tomography can be carried out even for nontrivial quantum evolution. The black-box approach of this work is an illustration of the fact that quantum mechanics is an operational theory, in which states and measurement outcomes can be predicted from physical observations without the mathematical abstraction of a Hilbert space. The model-agnostic nature of the RNN is therefore readily generalized to larger quantum systems. Such networks could excel at finding efficient state representations for larger systems, which could prove useful for real-time modeling, filtering, and parameter estimation. The robust, model-independent nature of prediction is a promising tool for the calibration of future quantum processors and will enable the characterization of imperfections outside of the scope of the usual approximation, such as correlated errors or non-Markovian noise, and may even be suited for identifying and quantifying effects initially unknown to the experimenter.

ACKNOWLEDGMENTS

We acknowledge M. Devoret, V. Ramasesh, J. Colless, and M. Blok for helpful discussions. L. S. M. acknowledge funding via National Science Foundation graduate student

fellowships. This research is supported in part by the U.S. Army Research Office under Grant No. W911NF-15-1-0496 and by the Air Force Office of Scientific Research under Grant No. FA9550-12-1-0378. E. F. acknowledges for funding via the LabEX ENS-ICFP Grants No. ANR-10-LABX-0010 and No. ANR-10-IDEX-0001-02 PSL*.

APPENDIX: EXPERIMENTAL SETUP AND SAMPLE

Our transmon qubit is fabricated on double-side-polished silicon, with a single double-angle-evaporated Al/AlO_x/Al Josephson junction. The internal dimensions of the aluminum 3D cavity are 81 mm × 51 mm × 20 mm. The qubit is positioned 23 mm from the edge of the cavity. The qubit is characterized by a charging energy $E_c/h = 220$ MHz and a 0 to 1 transition frequency $\omega_q = 2\pi \times 4.262$ GHz. The qubit coherence is characterized by an excited state lifetime of $T_1 = 60 \mu\text{s}$, echo time of $T_{\text{echo}} = 40 \mu\text{s}$, and Rabi decay time of $25 \mu\text{s}$. The lowest cavity mode used in the experiment has resonance frequency $\omega_c = 2\pi \times 6.666$ GHz, linewidth $\kappa = 2\pi \times 7.2$ MHz, and qubit dispersive frequency shift $\chi = 2\pi \times 0.18$ MHz. The cavity output is amplified using a lumped-element Josephson parametric amplifiers operated in the phase-sensitive mode, and the amplifier gain are set to 15 dB. The signal is further amplified with a cryogenic HEMT amplifier, model number LNF4-8. Owing to the high noise temperature of the HEMT, we added a Josephson traveling wave parametric amplifier immediately before it. At room temperature, the signal is further amplified and then mixed down to 105 MHz using a microwave mixer. Finally, this 105 MHz signal is read into a PC using two 1 gigasample per second analog-to-digital converters, where it is digitally demodulated down to dc, compressed, and sent to the Google Cloud platform for neural network training.

-
- [1] G. Chiribella, G. M. D'Ariano, and P. Perinotti, *Informational Derivation of Quantum Theory*, *Phys. Rev. A* **84**, 012311 (2011).
 - [2] K. Murch, S. Weber, C. Macklin, and I. Siddiqi, *Observing Single Quantum Trajectories of a Superconducting Quantum Bit*, *Nature (London)* **502**, 211 (2013).
 - [3] S. Weber, A. Chantasri, J. Dressel, A. N. Jordan, K. Murch, and I. Siddiqi, *Mapping the Optimal Route between Two Quantum States*, *Nature (London)* **511**, 570 (2014).
 - [4] S. Hacoen-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).
 - [5] 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).

- [6] M. Schuster and K. K. Paliwal, *Bidirectional Recurrent Neural Networks*, *IEEE Trans. Signal Process.* **45**, 2673 (1997).
- [7] T. Mikolov, M. Karafiát, L. Burget, J. Černocký, and S. Khudanpur, *Recurrent Neural Network Based Language Model*, in *Proceedings of the Eleventh Annual Conference of the International Speech Communication Association*, 2010, https://www.fit.vutbr.cz/research/groups/speech/publi/2010/mikolov_interspeech2010_IS100722.pdf.
- [8] A. Graves, A.-R. Mohamed, and G. Hinton, *Speech Recognition with Deep Recurrent Neural Networks*, in *Proceedings of the 2013 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)* (IEEE, New York, 2013), pp. 6645–6649.
- [9] L. Bouten, R. Van Handel, and M. R. James, *An Introduction to Quantum Filtering*, *SIAM J. Control Optimiz.* **46**, 2199 (2007).
- [10] G. Torlai, G. Mazzola, J. Carrasquilla, M. Troyer, R. Melko, and G. Carleo, *Neural-Network Quantum State Tomography*, *Nat. Phys.* **14**, 447 (2018).
- [11] L. Wang, *Discovering Phase Transitions with Unsupervised Learning*, *Phys. Rev. B* **94**, 195105 (2016).
- [12] J. Carrasquilla and R. G. Melko, *Machine Learning Phases of Matter*, *Nat. Phys.* **13**, 431 (2017).
- [13] E. P. Van Nieuwenburg, Y.-H. Liu, and S. D. Huber, *Learning Phase Transitions by Confusion*, *Nat. Phys.* **13**, 435 (2017).
- [14] G. Carleo and M. Troyer, *Solving the Quantum Many-Body Problem with Artificial Neural Networks*, *Science* **355**, 602 (2017).
- [15] T. Fösel, P. Tighineanu, T. Weiss, and F. Marquardt, *Reinforcement Learning with Neural Networks for Quantum Feedback*, *Phys. Rev. X* **8**, 031084 (2018).
- [16] G. Torlai and R. G. Melko, *Neural Decoder for Topological Codes*, *Phys. Rev. Lett.* **119**, 030501 (2017).
- [17] S. Krastanov and L. Jiang, *Deep Neural Network Probabilistic Decoder for Stabilizer Codes*, *Sci. Rep.* **7**, 11003 (2017).
- [18] P. Baireuther, T. E. O'Brien, B. Tarasinski, and C. W. Beenakker, *Machine-Learning-Assisted Correction of Correlated Qubit Errors in a Topological Code*, *Quantum* **2**, 48 (2018).
- [19] M. August and X. Ni, *Using Recurrent Neural Networks to Optimize Dynamical Decoupling for Quantum Memory*, *Phys. Rev. A* **95**, 012335 (2017).
- [20] I. Guevara and H. Wiseman, *Quantum State Smoothing*, *Phys. Rev. Lett.* **115**, 180407 (2015).
- [21] M. Tsang, *Time-Symmetric Quantum Theory of Smoothing*, *Phys. Rev. Lett.* **102**, 250403 (2009).
- [22] J. Koch, T. M. Yu, J. Gambetta, A. A. Houck, D. Schuster, J. Majer, A. Blais, M. H. Devoret, S. M. Girvin, and R. J. Schoelkopf, *Charge-Insensitive Qubit Design Derived from the Cooper Pair Box*, *Phys. Rev. A* **76**, 042319 (2007).
- [23] H. Paik, D. Schuster, L. S. Bishop, G. Kirchmair, G. Catelani, A. Sears, B. Johnson, M. Reagor, L. Frunzio, L. Glazman *et al.*, *Observation of High Coherence in Josephson Junction Qubits Measured in a Three-Dimensional Circuit QED Architecture*, *Phys. Rev. Lett.* **107**, 240501 (2011).
- [24] M. Hatridge, R. Vijay, D. H. Slichter, J. Clarke, and I. Siddiqi, *Dispersive Magnetometry with a Quantum Limited SQUID Parametric Amplifier*, *Phys. Rev. B* **83**, 134501 (2011).
- [25] 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).
- [26] F. A. Gers, J. Schmidhuber, and F. Cummins, *Learning to Forget: Continual Prediction with LSTM*, *Neural Comput.* **12**, 2451 (2000).
- [27] S. Hochreiter, Y. Bengio, P. Frasconi, J. Schmidhuber *et al.*, *Gradient Flow in Recurrent Nets: The Difficulty of Learning Long-Term Dependencies*, in *A Field Guide to Dynamical Recurrent Networks*, edited by S. C. Kremer and J. F. Kolen (IEEE Press, Piscataway, NJ, USA, 2001).
- [28] M. Abadi, P. Barham, J. Chen, Z. Chen, A. Davis, J. Dean, M. Devin, S. Ghemawat, G. Irving, M. Isard *et al.*, *Tensorflow: A System for Large-Scale Machine Learning*, in *Proceedings of the 12th USENIX Symposium on Operating Systems Design and Implementation (OSDI 16)* (2016), Vol. 16, pp. 265–283, <https://www.usenix.org/system/files/conference/osdi16/osdi16-abadi.pdf>.
- [29] D. P. Kingma and J. Ba, *Adam: A Method for Stochastic Optimization*, [arXiv:1412.6980](https://arxiv.org/abs/1412.6980).
- [30] N. Srivastava, G. Hinton, A. Krizhevsky, I. Sutskever, and R. Salakhutdinov, *Dropout: A Simple Way to Prevent Neural Networks from Overfitting*, *J. Mach. Learn. Res.* **15**, 1929 (2014).
- [31] K. Greff, R. K. Srivastava, J. Koutník, B. R. Steunebrink, and J. Schmidhuber, *LSTM: A Search Space Odyssey*, *IEEE Trans. Neural Networks Learn. Syst.* **28**, 2222 (2016).
- [32] P. Rouchon and J. F. Ralph, *Efficient Quantum Filtering for Quantum Feedback Control*, *Phys. Rev. A* **91**, 012118 (2015).
- [33] L. R. Rabiner, *A Tutorial on Hidden Markov Models and Selected Applications in Speech Recognition*, *Proc. IEEE* **77**, 257 (1989).
- [34] Y. Aharonov, P. G. Bergmann, and J. L. Lebowitz, *Time Symmetry in the Quantum Process of Measurement*, *Phys. Rev.* **134**, B1410 (1964).
- [35] S. Gammelmark, B. Julsgaard, and K. Mølmer, *Past Quantum States of a Monitored System*, *Phys. Rev. Lett.* **111**, 160401 (2013).
- [36] P. Campagne-Ibarcq, L. Bretheau, E. Flurin, A. Auffèves, F. Mallet, and B. Huard, *Observing Interferences between Past and Future Quantum States in Resonance Fluorescence*, *Phys. Rev. Lett.* **112**, 180402 (2014).
- [37] D. Tan, S. J. Weber, I. Siddiqi, K. Mølmer, and K. W. Murch, *Prediction and Retrodiction for a Continuously Monitored Superconducting Qubit*, *Phys. Rev. Lett.* **114**, 090403 (2015).
- [38] T. Rybarczyk, B. Peaudecerf, M. Penasa, S. Gerlich, B. Julsgaard, K. Mølmer, S. Gleyzes, M. Brune, J. M. Raimond, S. Haroche, and I. Dotsenko, *Forward-Backward*

- Analysis of the Photon-Number Evolution in a Cavity*, *Phys. Rev. A* **91**, 062116 (2015).
- [39] D. Tan, M. Naghiloo, K. Mölmer, and K. Murch, *Quantum Smoothing for Classical Mixtures*, *Phys. Rev. A* **94**, 050102 (R) (2016).
- [40] E. Magesan, J.M. Gambetta, A.D. Córcoles, and J.M. Chow, *Machine Learning for Discriminating Quantum Measurement Trajectories and Improving Readout*, *Phys. Rev. Lett.* **114**, 200501 (2015).
- [41] A.S. Holevo, *Bounds for the Quantity of Information Transmitted by a Quantum Communication Channel*, *Prob. Peredachi Inf.* **9**, 3 (1973).
- [42] P. Six, P. Campagne-Ibarcq, I. Dotsenko, A. Sarlette, B. Huard, and P. Rouchon, *Quantum State Tomography with Noninstantaneous Measurements, Imperfections, and Decoherence*, *Phys. Rev. A* **93**, 012109 (2016).