

Experimental quantum reservoir computing with a circuit-quantum-electrodynamics system

B. Carles¹, J. Dudas¹, L. Balembois², J. Grollier¹ and D. Marković^{1,*}

¹*Laboratoire Albert Fert, CNRS, Thales, Université Paris-Saclay, 91767 Palaiseau, France*

²*SPEC, CEA, Gif-sur-Yvette, France*

 (Received 26 June 2025; revised 20 November 2025; accepted 2 February 2026; published 4 May 2026)

Quantum reservoir computing is a machine-learning framework that offers ease of training compared to other quantum neural networks, as it does not rely on gradient-based optimization. Learning is performed in a single step on the output features measured from the quantum system. Various implementations of quantum reservoir computing have been explored in simulations, with different measured features. Although simulations have shown that quantum reservoirs present advantages in performance compared to classical reservoirs, experimental implementations have remained scarce. This is due to the challenge of obtaining a large number of output features that are nonlinear transformations of the input data. In this work, we show that even with a circuit quantum electrodynamics system as simple as a single transmon coupled to a readout resonator, we can implement a proof-of-concept realization of quantum reservoir computing. We obtain a large number of nonlinear features from a single physical system by encoding the input data in the amplitude of a coherent drive and measuring the cavity state in the Fock basis. We demonstrate classification of two classical tasks with significantly smaller hardware resources and fewer measured features compared to classical neural networks. Our experimental results are supported by numerical simulations that show additional Kerr nonlinearity is beneficial to reservoir performance. Our work demonstrates a hardware-efficient quantum neural-network implementation that can be further scaled up and generalized to other quantum machine-learning models.

DOI: [10.1103/fmm7-3qm9](https://doi.org/10.1103/fmm7-3qm9)

I. INTRODUCTION

Quantum reservoir computing (QRC) was introduced as an approach to quantum machine learning that circumvents the need to train quantum systems directly. This is particularly appealing because gradient-based optimization methods that work so well for classical artificial neural networks are difficult to apply to quantum systems: estimating gradients is experimentally demanding, and the optimization landscape often suffers from barren plateaus that hinder convergence [1]. Quantum reservoir computing has been explored in simulations using a variety of systems, including qubit-based reservoirs [2,3] and bosonic modes [4–6]. These studies have shown that QRC can solve nontrivial tasks with fewer output neurons than classical reservoirs, highlighting both the computational advantage provided by quantum coherence and the potential for processing

quantum input data [3]. Experimental realizations of quantum reservoir computing, however, remain scarce. Existing implementations include demonstrations on qubit-based quantum processors available through online platforms [7–9], quantum extreme learning machines with nuclear spins [10] and with photon orbital-angular momentum [11,12] and, more recently, an analog realization using a bosonic mode coupled to a qubit [13]. Here, we report on a hardware-efficient implementation of quantum reservoir computing using a single bosonic mode hosted by a superconducting resonator, coupled to an ancilla qubit for readout. In our scheme, the input data are encoded in the amplitude of the cavity displacement, while the reservoir outputs are given by the Fock-state occupation probabilities. This approach provides a large number of nonlinear features directly from the quantum dynamics. Of note, the nonlinearity arises from the measurement process itself rather than from the input encoding, a key property for future extensions toward learning on quantum data.

II. CIRCUIT AND IMPLEMENTATION

We implement quantum reservoir computing on a superconducting circuit composed of a tantalum coplanar waveguide resonator, capacitively coupled to a transmon

*Contact author: danijela.markovic@cnrs.fr

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




FIG. 1. Quantum reservoir computing with a superconducting circuit. (a) Input data are encoded in the amplitude of the resonant drive (green) of the quantum system acting as a reservoir (blue). Output features (purple circles) are obtained by measuring the occupation probabilities of Fock states $|0\rangle$ to $|n\rangle$ at different times t_1 and t_2 , yielding a total of $2n$ output features. Features are classified by training a layer of linear weights (dashed lines). (b) Schematic and (c) design of the quantum circuit used to implement the quantum reservoir. Tantalum coplanar waveguide resonator (blue) is capacitively coupled to a transmon qubit (purple) and to a transmission line through a Purcell filter (green).

qubit [Figs. 1(b) and 1(c)]. We use the fundamental $\lambda/2$ mode of the resonator as a cavity, at a frequency $\omega_c = 2\pi \times 7.617$ GHz. We use the transmon qubit as ancilla, to measure the occupation probabilities of the cavity Fock states that serve the function of output neurons [Fig. 1(a)]. The cavity-qubit system can be described by the Hamiltonian

$$\begin{aligned} \frac{\hat{H}}{\hbar} = & \omega_c \hat{a}^\dagger \hat{a} + \omega_q \frac{\hat{\sigma}_Z}{2} - \frac{\chi}{2} \hat{a}^\dagger \hat{a} \hat{\sigma}_Z + K_{cc} (\hat{a}^\dagger \hat{a})^2 \\ & + K_{cq} (\hat{a}^\dagger \hat{a})^2 \frac{\hat{\sigma}_Z}{2}, \end{aligned} \quad (1)$$

where $\omega_q = 2\pi \times 6.210$ GHz is the bare qubit frequency, $\chi = 2\pi \times 22.29$ MHz is the dispersive coupling rate, $K_{cc} = -2\pi \times 300$ kHz is the cavity self-Kerr coefficient and $K_{cq} \approx -2\pi \times 0.44$ MHz is a photon-number-dependent correction to the dispersive shift. The dispersive coupling makes the qubit resonance frequency depend on the number of photons $n = \langle \hat{a}^\dagger \hat{a} \rangle$ in the cavity.




FIG. 2. Fock-state probabilities as neurons. (a) Two-tone spectroscopy of the cavity. The phase shift of the reflection coefficient at the cavity resonance frequency, as a function of the qubit drive frequency and the cavity drive amplitude $\alpha_{in}(t)$. Different resonances from right to left correspond to the qubit frequency dressed by 0, 1, 2, 3, and 4 photons in the cavity. Yellow horizontal lines indicate the amplitude range used for data encoding. (b) Top: pulse sequence used to measure the photon number in the cavity. A 200-ns displacement pulse D_α at cavity frequency ω_c is followed by a 200-ns π_n pulse conditioned on n photons in the cavity, and a high power readout (HPR) after the 1- μ s waiting time. Bottom: measured occupation probabilities (dots) P_0 to P_4 for Fock states $|0\rangle$ to $|4\rangle$ as a function of the cavity displacement amplitude α_{in} . Each point is averaged 10 000 times. Lines correspond to the probabilities obtained from the simulations using the Lindblad master equation, with dissipation rates and Kerr effect as free parameters.

The cavity drive Hamiltonian is

$$\hat{H}_{\text{drive}} = \hbar \sqrt{\kappa_{\text{ext}}} \alpha_{in}(t) (\hat{a} + \hat{a}^\dagger), \quad (2)$$

where κ_{ext} is the cavity coupling to the transmission line and $\alpha_{in}(t)$ is the amplitude of the classical cavity drive. In the two-tone spectroscopy shown in Fig. 2(a), we observe a set of equally spaced drive frequencies that lead to a significant shift of the cavity resonance. These drive frequencies, separated by χ , correspond to resonance frequencies $\omega_q^n = \omega_q - n\chi$ of the qubit dressed by n photons in the cavity. As is commonly done in circuit QED [14], we use these photon-number-dependent resonances to measure the number of photons in the cavity, and to infer the Fock-state occupation probabilities from the measurement statistics. The time sequence used to measure the occupation probabilities is shown in the top of Fig. 2(b). A 200-ns cavity displacement pulse D_α is applied at cavity resonance associated with the transmon being in the ground state, followed by a photon-number-conditioned π_n pulse applied to the qubit at frequency ω_q^n . The state of the qubit is read through the same resonator, using high-power readout [15]. For this reason, we perform the readout 1 μ s after the π_n pulse, letting enough time for the resonator to empty, such that the readout pulse is not affected by the photons residual

from the input data. In future experiments, this delay could be reduced by employing a dedicated readout resonator coupled to the qubit. Furthermore, the protocol could be made more efficient by using the multiplexed readout [16,17] and deep neural-network discrimination [18].

III. CHARACTERIZATION OF NEURAL NONLINEARITIES

Fock-state populations of a coherent state $|\alpha\rangle$ are given by Poissonian distributions, $P_n(\alpha) = e^{-|\alpha|^2} \frac{|\alpha|^{2n}}{n!}$, which have a nonlinear dependence on the cavity coherent state amplitude α , that can be exploited as a neural activation function. In our circuit, the measured distributions deviate from Poisson statistics as a result of the Kerr effect originating from the nonlinearity inherited by the cavity from the qubit. We simulate the probabilities using the Lindblad master equation, with the dissipation rates and cavity self-Kerr coefficient as free parameters determined by fitting to the measured data [Fig. 2(b)]. The extracted values, $\kappa_{\text{tot}} = \kappa_{\text{ext}} + \kappa_{\text{int}} = 2\pi \times (560 \pm 60)$ kHz and $K_{cc} = 2\pi \times (-300 \pm 50)$ kHz, are consistent with those obtained from reflection measurements.

Due to the self-Kerr effect, the cavity mode frequency shifts at high input amplitudes corresponding to large data values. Within the encoding range $\alpha_{\text{in}} \in [1, 10.4]\sqrt{\text{MHz}}$, the cavity resonance shifts by $\Delta f_c = 1.7$ MHz, which largely exceeds the cavity linewidth κ_{tot} . To account for this shift, we use short input pulses of duration $\Delta t_{\text{in}} = 200$ ns, giving a spectral width $\Delta t_{\text{in}}^{-1} = 5$ MHz $> \Delta f_c$. Furthermore, a higher-order correction, corresponding to a photon-number-dependent dispersive shift K_{cq} , reduces the efficiency of the conditional π_n pulse and, consequently, the overall readout efficiency.

IV. RESULTS: LEARNING TEMPORAL TASKS

A. Sine and square waveform classification

We first deploy our quantum reservoir computing scheme on the task of sine and square waveform classification. This task consists in sending one by one, points that belong either to a sine or to a square period, and assigning them the correct label, as shown in Fig. 3(a). This task probes both the nonlinearity and memory of the underlying physical system: a fully linear network can only apply a single amplitude threshold, achieving at best 68.75% accuracy, while classical recurrent neural networks combining nonlinearity and memory can reach 100% accuracy with 25 neurons [19].

We have built a dataset of 400 random sine and square periods, each discretized into eight points. Half of the data is used for training and the remaining half for testing. The prediction

$$Y = FW \quad (3)$$

is obtained by multiplying the feature matrix F by the weight matrix W . The feature matrix is composed of the occupation probabilities of states $|0\rangle$ to $|4\rangle$ measured at times $t_i = i \times 50$ ns, $i \in [1, 4]$ after the beginning of the pulse, and a bias term. Each probability is reconstructed from 4000 measurements.

After each measurement, the reservoir state is reinitialized by waiting 60 μ s for the qubit to relax to its ground state before reinjecting the same input x_n preceded by the input x_{n-1} , because this task requires the memory of at least one previous input. This procedure yields a total of $N_F = 20$ output features, corresponding to five probabilities at four different times. The weight matrix W is computed on the training data using the ridge regression, by minimizing the cost function

$$J(W) = |Y - FW|^2 + \beta|W|^2. \quad (4)$$

The second term is a penalty on large weights, controlled by the regularization parameter β . With 20 measured features, yielding 21 output neurons including the bias, we find that the prediction on the test set matches the target with 99.8% accuracy. Interestingly, even when we retain only the eight most informative features, $P_0^t, P_0^t, P_0^t, P_1^t, P_2^t, P_2^t, P_3^t, P_4^t$, we still achieve 99.5% accuracy, as shown in Fig. 3(c). This is a significantly smaller number of features that need to be measured compared to classical reservoir computing, where at least 25 classical features are required to achieve comparable performance, and is consistent with previous simulations [19]. We can notice that the most informative features correspond to a mix of different states measured at different times. In a practical implementation, the most informative features could be identified during the training phase, such that during inference only they are measured.

In practice, the extraction of these features is subject to experimental imperfections that distort the measured occupation probabilities. Specifically, we measure the probability of ionizing the qubit after a π_n pulse followed by a high-power readout pulse. Under ideal conditions, this probability would be equal to the underlying Fock-state probability P_n . However, due to thermal noise, finite efficiency of the π_n pulse, finite ionization probability and qubit decay during the waiting time between the π_n pulse and the readout, the measured switching probability $p_s(n)$ becomes an affine function of the true probability, $p_s(n) = aP_n + b$. These effects modify the output probabilities but do not degrade the performance of the quantum reservoir, since the learned weights can compensate for such distortions. Measurement imperfections nevertheless tend to concentrate $p_s(n)$ toward 0.5, increasing the number of measurement shots required to reach a given statistical precision.

Finally, beyond such technical imperfections, we analyze the impact of decoherence on the learning




FIG. 3. Sine and square waveform classification task. (a) Input data are obtained from a random series of 400 sine and square periods, discretized into eight points per period. Only four periods are shown for clarity. (b) Top: pulse sequence consists of two displacement pulses $D_{\alpha_{i-1}}$ and D_{α_i} , whose amplitudes encode two consecutive data points, followed by a π_n pulse and high-power readout. To increase the number of extracted features, the occupation probabilities are sampled at four times t_1 to t_4 during the second displacement pulse. Bottom: cavity drive $\text{Re}(\alpha_{\text{in}})$ whose amplitude encodes the input data. (c) Prediction of the quantum reservoir with the eight most informative features among probabilities of states $|0\rangle$ to $|4\rangle$ at times $t_i = i \times 50$ ns, $i \in [1, 4]$. (d) Experimental accuracy (purple crosses) as a function of the number of measured states N , with four time samples per state, giving $4N$ features. For each N , the plot shows the best accuracy obtained with the N most informative features. The solid lines connect simulated data points (dots) corresponding to different qubit dephasing rates κ_ϕ . (e) Experimental accuracy with five measured features (states from $|0\rangle$ to $|4\rangle$) as a function of the number of time sampling steps, with sampling times indicated in brackets. (f) Accuracy as a function of the number of measurement shots and (g) as a function of the encoding range.

performance. To pinpoint the role of quantum coherence in learning this task [20], we simulate the system for different values of the dephasing rate κ_ϕ . We observe that the accuracy decreases significantly as dephasing increases, even when a larger number of Fock states is measured [Fig. 3(d)]. Figures 3(e)–3(g) show the dependence of the experimental reservoir performance on the number of time-sampling steps, the finite number of measurement shots [21], and the amplitude encoding range. Interestingly, we find that using a larger encoding range is advantageous, despite the fact that at higher drive amplitudes the cavity population saturates due to the self-Kerr effect, and the π_n pulses become less efficient because of the cross-Kerr effect. This behavior will be further analyzed in simulations described in Sec. V.

B. Mackey-Glass chaotic time-series prediction

We implement a second benchmark task, which is Mackey-Glass chaotic time-series prediction [22]. This

task consists in forecasting future values of a time series based on its past values. The time series is generated as a solution of the Mackey-Glass nonlinear differential equation

$$\frac{\partial x(t)}{\partial t} = \frac{\beta x(t - \tau)}{1 + x(t - \tau)^{10}} - \gamma x(t), \quad (5)$$

and is chaotic for $\beta = 0.2$, $\gamma = 0.1$, and $\tau = 17$. The whole series contains 2000 data points. We use half of the dataset for training and the remaining half for testing. For each value of the delay d , the system is trained to try to output a time-delayed version of the Mackey-Glass time series. The target is thus $\tilde{y}_k = x_{k+d}$, as shown in Fig. 4(b). The feature vector is composed of five occupation probabilities of states $|0\rangle$ to $|4\rangle$ measured 100 ns after the beginning of the input encoding pulse. For each measurement of the feature vector for the input k , we reinject 20 preceding inputs $[x_{k-19}, \dots, x_k]$, each for 100 ns. We evaluate performance using the normalized root-mean-square




FIG. 4. Mackey-Glass chaotic time-series prediction. (a) Pulse sequence used for the task. (b) Part of the input data (green) and corresponding target for a delay $d = 25$ (gold). Quantum reservoir predictions for delays $d = 1$ and $d = 25$ are shown, respectively, in pink and purple. (c) Normalized root-mean-square error (NRMSE) as a function of delay d for the experimental quantum reservoir for the lag parameter $\tau = 17$. (d) NRMSE as a function of the delay d for a simulated quantum reservoir for different parameters τ .

error,

$$\epsilon = \frac{1}{\Delta \tilde{y}} \sqrt{\frac{\sum_k (y_k - \tilde{y}_k)^2}{N}}, \quad (6)$$

where the normalization is ensured by the division by the range of the target values $\Delta \tilde{y} = \max(\tilde{y}) - \min(\tilde{y})$.

Figure 4(c) shows the NRMSE as a function of the prediction delay d for the lag parameter $\tau = 17$, value for which the Mackey-Glass series becomes quasichaotic. We observe that the NRMSE exhibits minima for delays d corresponding to multiples of the quasioscillation period, for which the target signal is highly correlated with the input. To further illustrate this effect, we show in Fig. 4(d) simulated NRMSE values for Mackey-Glass time series with different delay parameters τ , ranging from smaller (periodic) to larger (chaotic) regimes. For periodic inputs, clear oscillations in the NRMSE are visible, whereas for highly chaotic inputs, the NRMSE rapidly converges to an average value.

V. SIMULATIONS OF THE IMPACT OF THE KERR EFFECT

Having experimentally demonstrated the classification of sine and square waveforms, as well as the prediction of a chaotic time series using a quantum reservoir, we now turn to simulations to explore the physical origin of the reservoir's computational capabilities. In particular, we investigate how the Kerr nonlinearity—which plays a central role in shaping the system's dynamics—affects performance.

The effect of different types and magnitudes of nonlinearities on the performance of quantum neural networks remains an open question [23,24]. In optical quantum neural networks, some studies have suggested that weaker Kerr nonlinearities lead to degraded performance [25], while in the context of bosonic quantum neural networks, it has been shown that whereas simple classical tasks such as XOR gate learning can be solved with only the encoding nonlinearity, quantum tasks require an additional Kerr nonlinearity, which also improves robustness to errors and noise [26].

To investigate the influence of Kerr nonlinearity in our setup, we simulate the classification of sine and square waveforms using a Kerr oscillator with varying Kerr coefficients K_{cc} . Since the input drive is applied at a fixed frequency for all drive amplitudes, its coupling to the resonator becomes less efficient when the Kerr coefficient is large. This significantly affects the mode population. Figure 5(a) shows the average photon number in the mode as a function of the drive detuning from the cavity resonance at low power, $\Delta = \omega_c - \omega$, and the drive amplitude α_{in} , after two 200-ns input pulses. We observe that for large Kerr coefficients K_{cc} and zero detuning, the mean photon number remains below 2. Indeed, in this case, the drive is too far from the effective resonance to efficiently populate the cavity. For larger detunings, the mean photon number is a nonmonotonic function of the drive amplitude: it initially increases, reaching a maximum at an intermediate amplitude where the drive frequency is closest to resonance, and then decreases again at higher amplitudes as the system




FIG. 5. Impact of the Kerr effect on performance. (a) Simulated mean photon number $\langle n \rangle$ in the cavity after a 400-ns drive as a function of detuning and input amplitude α_{in} , for different Kerr rates K_{cc} . Contour lines indicate regions of equal mean photon number. (b) NRMSE for sine-square classification using 20 neurons, shown as a function of Kerr coefficient and detuning, for various input encoding ranges $[\alpha_{\text{in}}^{\text{min}}, \alpha_{\text{in}}^{\text{max}}]$. (c) Simulated NRMSE as a function of the Kerr coefficient K_{cc} for $\Delta = 0$ and the encoding range $[0.7\alpha_0, 1.3\alpha_0]$, where α_0 is the drive amplitude required to reach $\langle n \rangle \in \{1, 2, 3\}$ after a 400-ns drive [see panel (a)]. Because multiple combinations of detuning and amplitude can yield the same mean photon number, the NRMSE values for different amplitudes are shown with distinct line styles.

moves away from resonance. This implies that for large Kerr coefficients, it will potentially be interesting to apply the input data at a frequency detuned by $\Delta \neq 0$ from the resonance.

In the following, we study the performance on the sine-square task as a function of Kerr coefficient K_{cc} . In order to separate the impact of the additional nonlinearity from the spurious effects such as the photon number in the cavity related to the frequency shift, we perform the study for different drive detunings and also for different encoding ranges $[\alpha_{\text{min}}, \alpha_{\text{max}}]$ as shown in Fig. 5(b). We observe that the lowest classification errors are achieved for the largest encoding ranges and for high values of the Kerr coefficient.

To isolate the contribution of the Kerr effect from that of the amplitude encoding range, we proceed as follows: for a given Kerr coefficient and a fixed target mean photon number $\langle n \rangle \in \{1, 2, 3\}$, we select all combinations of $\{\Delta, \alpha_0\}$ such that the lowest input amplitude results in $\langle n \rangle$

photons in the cavity. We then define the encoding range as $[0.7\alpha_0, 1.3\alpha_0]$. As shown in Fig. 5(c), we observe that the lowest errors are obtained for large Kerr coefficients. It is worthwhile to note that due to the non-monotonic behavior of the photon number as a function of input amplitude, for a given Kerr coefficient, different drive amplitudes can yield the same average photon number, as shown in Fig. 5(a). This explains the presence of multiple points with the same Kerr coefficient but different amplitudes in Fig. 5(c). Lower errors are generally associated with higher drive amplitudes. However, in practice, the maximum amplitude is limited by the onset of parametric oscillation [27].

VI. CONCLUSION AND OUTLOOK

We have experimentally implemented quantum reservoir computing using a hardware-efficient platform composed of a single bosonic mode coupled to a qubit. By encoding input data in the cavity-mode displacement amplitudes and performing readout in the Fock basis, we extract a large number of nonlinear output features. This is essential in reservoir computing, where only the output weights are trained, and therefore a rich set of features is required. Additional features can alternatively be generated through time multiplexing [2], by weak continuous monitoring of multiple observables [28], by increasing the number of physical components in the system [7, 29], or by computing statistical moments of the measured data [13].

Nonlinear output features are necessary for successful information processing. In quantum bosonic systems, nonlinear transformations may arise from the encoding of input data [5], from intrinsic Kerr nonlinearity [4, 6], or from the measurement process [13, 19]. However, for applications involving quantum input data [3, 30–32], where the reservoir must interface seamlessly with a quantum system, nonlinearity cannot rely solely on the encoding stage. While Kerr nonlinearity is essential for bosonic reservoirs that use continuous variables as output features, in implementations such as ours that use discrete-variable outputs, the Kerr effect adds to the measurement-induced nonlinearity. To investigate its influence, we performed simulations isolating the effect of Kerr nonlinearity on classification performance. Our results show that Kerr nonlinearity is beneficial, especially in regimes where the input amplitudes interact nontrivially with the detuned resonator dynamics. These findings identify Kerr-induced nonlinearity as a useful computational resource, whose role in temporal and quantum tasks will be explored in future work. Looking forward, this scheme can be scaled up by capacitively or parametrically coupling multiple resonators, each with a dedicated ancilla qubit. In such scaled architectures, the robustness of trained readouts to slow parameter drifts and environmental nonstationarities will become an increasingly important consideration. In the

present work, all training and testing data are acquired within a single continuous measurement session, during which the system dynamics remain stationary and the trained readout remains valid. While this justifies the performance reported here, the long-term stability of trained reservoirs has not been systematically investigated. Quantifying drift timescales, sensitivity to small fluctuations, and the associated retraining requirements will be crucial for scalable circuit QED-based reservoir computing and constitutes a key direction for future work.

ACKNOWLEDGMENTS

The authors thank E. Flurin (SPEC, CEA, Gif-sur-Yvette, France) for his help in the sample fabrication and advice on performing the circuit measurements. This research was supported by the Paris Ile-de-France Region in the framework of DIM SIRTEQ and by European Union (ERC, qDynnet, 101 076 898).

DATA AVAILABILITY

The data that support the findings of this article are openly available [34], embargo periods may apply.

APPENDIX A: DEVICE AND MEASUREMENT SETUP

1. Device

The device is a single microwave photon detector shown in Fig. 6 fabricated by L. Balembois and E. Flurin in SPEC, CEA [33]. In the present work, we use the waste resonator (green) as the cavity. The relevant circuit parameters are summarized in Table I.

2. Measurement setup

The experimental setup is shown in Fig. 6. The cavity and qubit drives are generated by an advanced quantum control and processing unit, the OPX+ by Quantum Machines, which comprises an arbitrary waveform generator (AWG) and a fast acquisition card. Pulses are modulated at 233.88 MHz for the cavity and at 241.31, 219.3, 198.0, 177.1, and 157.4 MHz for the qubit, with a sampling rate of 1 GHz. *I-Q* mixers are used to combine these rf pulses with local-oscillator tones generated by an Anapico APMS20G at 7.383 12 and 5.969 GHz.

The reflected signal is amplified by a HEMT amplifier at 4 K and by three 15-dB amplifiers at room temperature. The signal is then down-converted using *I-Q* mixers and digitized by the OPX+. dc blocks are added on the input




FIG. 6. Schematic of the experimental setup at low and room temperature. The sample is placed on the base stage of a dilution cryostat at 11 mK, inside a cryoperm electromagnetic shield. The sample design is adapted from Ref. [33]. In our current setup, the flux noise was too strong, making the second resonator (yellow) unusable in practice.

TABLE I. Circuit parameters.

Circuit parameter	Symbol	Value
Cavity frequency	$\omega_c/2\pi$	7.617 GHz
Qubit frequency	$\omega_q/2\pi$	6.210 31 GHz
Cavity-qubit cross Kerr rate	$\chi/2\pi$	22.29 MHz
Qubit decay time	T_1	8.01 μ s
Cavity decay time	T_c	0.93 μ s

lines to maintain a 50- Ω impedance match, as well as after amplification to improve digitization. The OPX+ further amplifies the output signal by 20 dB, digitizes it at a 1-GHz sampling rate, and performs the demodulation.

APPENDIX B: CALIBRATION

The cavity frequency ω_c and decay time T_c are determined by fitting the cavity reflection coefficient. The occupation probabilities of the cavity are measured using qubit π pulses conditioned on the cavity Fock state $|n\rangle$, denoted $\pi_{|n\rangle}$. The $\pi_{|0\rangle}$ pulse is implemented using a Gaussian pulse with a duration of 200 ns, corresponding to approximately $5/\chi$. Its frequency and amplitude are determined by fitting Rabi fringe measurements. All the other $\pi_{|n\rangle}$ pulses have the same shape, amplitude, and duration, while their frequencies are extracted from the two-tone spectroscopy shown in Fig. 7. Specifically, the cavity output field amplitude α_{out} is measured for different cavity displacements α_{in} as a function of the frequency of a subsequent 200-ns qubit pulse. Dressed qubit resonance frequencies are extracted from the fitted minima of the cavity output field amplitude. The distance between different resonances corresponds to the cavity-qubit cross-Kerr rate χ .

Qubit state measurement is performed using the high-power readout (HPR) as described in Ref. [15]. Figure 8




FIG. 7. Amplitude of the cavity output field at the end of the pulse sequence as a function of the qubit drive frequency for different color-coded cavity displacement amplitudes. The pulse sequence is composed of a cavity displacement pulse, followed by a 200-ns qubit pulse and a high cavity readout pulse. Full lines are fits centered around ω_q^n for the amplitude that yields n photons in the cavity with highest probability.




FIG. 8. Integrated cavity output signal on 10 000 experiments represented in the quadrature space without the $\pi_{|0\rangle}$ pulse (top) and after a conditional $\pi_{|0\rangle}$ pulse on the qubit (bottom). The two clusters correspond to cavity states dressed by a qubit in the excited state (left cluster) or ground state (right cluster). The probability to ionize the qubit and thus to switch the cavity state corresponds to the proportion of points below the threshold $\text{Re}(\alpha_{\text{out}}) = 0.15$ (dashed black line).

illustrates this technique by showing that the cavity response to a high-power drive at its bare frequency results in two distinct states in the quadrature space. These two states correspond to the cavity dressed by an ionized or nonionized qubit and can be separated by applying a threshold at $\text{Re}(\alpha_{\text{out}}) = 0.15$ (dashed line). Indeed, high-power cavity drive can lead to qubit ionization and thus a switch of the cavity resonant frequency. Since the cavity bare frequency corresponds to the resonance frequency of the cavity dressed by the ionized qubit, probing at this frequency results in a switch from a cavity dark state to a cavity bright state when the qubit becomes ionized. The readout amplitude is a key parameter of this measurement: if it is too low, the qubit will not be ionized even if it is initially in the excited state, whereas if it is too high, the qubit will be ionized even when it is in the ground state. The readout amplitude is therefore chosen to maximize the contrast between the switching probabilities with and without a $\pi_{|0\rangle}$ pulse.

Furthermore, since the same cavity is used both for state preparation and readout, it is necessary to wait for cavity depletion before applying the HPR pulse. Otherwise, the effective readout amplitude would be increased by residual photons in the cavity, leading to unwanted qubit ionization and an overestimation of the switching probability. To identify the necessary delay time, we measure the occupation probabilities of states $|0\rangle$ and $|4\rangle$ for three different delays (0, 1, and 3 μ s), shown in Fig. 9. The occupation probabilities P_n are obtained by correcting for imperfections in the HPR measurement to extract p_e , and then




FIG. 9. Occupation probabilities of the cavity states $|0\rangle$ (circles) and $|4\rangle$ (stars) as a function of the cavity-displacement amplitude, for different color-coded delays between the conditional $\pi_{|n\rangle}$ pulse and the high-power readout (HPR) pulse. Probabilities saturate for delays longer than $1\ \mu\text{s}$ and are overestimated for shorter delays. Inset: pulse sequence associated with the measurement. The cavity displacement pulse (green) is sent at the cavity resonant frequency for the qubit in the ground state, $f_{\text{cav}}^{(g)}$. The π pulse conditioned on the cavity Fock state $|n\rangle$ is applied to the qubit (dark purple). Finally, the cavity is driven by the high-power readout pulse at the cavity resonant frequency for the ionized qubit state, $f_{\text{cav}}^{(\text{ion})}$.

multiplying by e^{T/T_q} , where T_q is the qubit lifetime, to compensate for qubit decay during the delay. We observe that the probabilities saturate for delays longer than $1\ \mu\text{s}$ and thus adopt the $1\text{-}\mu\text{s}$ delay between the $\pi_{|n\rangle}$ pulse and the readout for all the measurements. The reflected signal is recorded and demodulated to extract its envelope α_{out} , and a threshold $\text{Re}(\alpha_{\text{out}}) = 0.15\ \text{mV}$ is applied to determine the output state.

[1] M. Larocca, S. Thanasilp, S. Wang, K. Sharma, J. Biamente, P. J. Coles, L. Cincio, J. R. McClean, Z. Holmes, and M. Cerezo, Barren plateaus in variational quantum computing, *Nat. Rev. Phys.* **7**, 174 (2025).
 [2] K. Fujii and K. Nakajima, Harnessing disordered-ensemble quantum dynamics for machine learning, *Phys. Rev. Appl.* **8**, 024030 (2017).
 [3] S. Ghosh, A. Opala, M. Matuszewski, T. Paterek, and T. C. H. Liew, Quantum reservoir processing, *Npj Quantum Inf.* **5**, 35 (2019).
 [4] L. C. G. Govia, G. J. Ribeill, G. E. Rowlands, H. K. Krovi, and T. A. Ohki, Quantum reservoir computing with a single nonlinear oscillator, *Phys. Rev. Res.* **3**, 013077 (2021).
 [5] J. Nokkala, R. Martínez-Peña, G. L. Giorgi, V. Parigi, M. C. Soriano, and R. Zambrini, Gaussian states of continuous-variable quantum systems provide universal and versatile reservoir computing, *Commun. Phys.* **4**, 53 (2021).
 [6] G. Angelatos, S. A. Khan, and H. E. Türeci, Reservoir computing approach to quantum state measurement, *Phys. Rev. X* **11**, 041062 (2021).

[7] F. Hu, S. A. Khan, N. T. Bronn, G. Angelatos, G. E. Rowlands, G. J. Ribeill, and H. E. Türeci, Overcoming the coherence time barrier in quantum machine learning on temporal data, *Nat. Commun.* **15**, 7491 (2024).
 [8] Y. Suzuki, Q. Gao, K. C. Pradel, K. Yasuoka, and N. Yamamoto, Natural quantum reservoir computing for temporal information processing, *Sci. Rep.* **12**, 1353 (2022).
 [9] J. Chen, H. I. Nurdin, and N. Yamamoto, Temporal information processing on noisy quantum computers, *Phys. Rev. Appl.* **14**, 024065 (2020).
 [10] M. Negoro, K. Mitarai, K. Fujii, K. Nakajima, and M. Kitagawa, Machine learning with controllable quantum dynamics of a nuclear spin ensemble in a solid, *arXiv:1806.10910*.
 [11] A. Suprano, D. Zia, L. Innocenti, S. Lorenzo, V. Cimini, T. Giordani, I. Palmisano, E. Polino, N. Spagnolo, F. Sciarrino, G. M. Palma, A. Ferraro, and M. Paternostro, Experimental property reconstruction in a photonic quantum extreme learning machine, *Phys. Rev. Lett.* **132**, 160802 (2024).
 [12] D. Zia, L. Innocenti, G. Minati, S. Lorenzo, A. Suprano, R. D. Bartolo, N. Spagnolo, T. Giordani, V. Cimini, G. M. Palma, A. Ferraro, F. Sciarrino, and M. Paternostro, Quantum reservoir computing for photonic entanglement witnessing, *Sci. Adv.* **11**, eady7987 (2025).
 [13] A. Senanian, S. Prabhu, V. Kremenetski, S. Roy, Y. Cao, J. Kline, T. Onodera, L. G. Wright, X. Wu, V. Fatemi, and P. L. McMahon, Microwave signal processing using an analog quantum reservoir computer, *Nat. Commun.* **15**, 7490 (2024).
 [14] D. I. Schuster, A. A. Houck, J. A. Schreier, A. Wallraff, J. M. Gambetta, A. Blais, L. Frunzio, J. Majer, B. Johnson, M. H. Devoret, S. M. Girvin, and R. J. Schoelkopf, Resolving photon number states in a superconducting circuit, *Nature* **445**, 515 (2007).
 [15] M. D. Reed, L. DiCarlo, B. R. Johnson, L. Sun, D. I. Schuster, L. Frunzio, and R. J. Schoelkopf, High-fidelity readout in circuit quantum electrodynamics using the Jaynes-Cummings nonlinearity, *Phys. Rev. Lett.* **105**, 173601 (2010).
 [16] J. Heinsoo, C. K. Andersen, A. Remm, S. Krinner, T. Walter, Y. Salathé, S. Gasparinetti, J.-C. Besse, A. Potočnik, A. Wallraff, and C. Eichler, Rapid high-fidelity multiplexed readout of superconducting qubits, *Phys. Rev. Appl.* **10**, 034040 (2018).
 [17] A. Essig, Q. Ficheux, T. Peronnin, N. Cottet, R. Lescanne, A. Sarlette, P. Rouchon, Z. Leghtas, and B. Huard, Multiplexed photon number measurement, *Phys. Rev. X* **11**, 031045 (2021).
 [18] B. Lienhard, A. Vepsäläinen, Luke C. G. Govia, C. R. Hoffer, J. Y. Qiu, D. Risté, M. Ware, D. Kim, R. Winik, A. Melville, B. Niedzielski, J. Yoder, G. J. Ribeill, T. A. Ohki, H. K. Krovi, T. P. Orlando, S. Gustavsson, and W. D. Oliver, Deep-neural-network discrimination of multiplexed superconducting-qubit states, *Phys. Rev. Appl.* **17**, 014024 (2022).
 [19] J. Dudas, B. Carles, E. Plouet, F. A. Mizrahi, J. Grollier, and D. Marković, Quantum reservoir computing implementation on coherently coupled quantum oscillators, *Npj Quantum Inf.* **9**, 64 (2023).

- [20] A. Palacios, R. Martínez-Peña, M. C. Soriano, G. L. Giorgi, and R. Zambrini, Role of coherence in many-body quantum reservoir computing, *Commun. Phys.* **7**, 369 (2024).
- [21] S. A. Khan, F. Hu, G. Angelatos, and H. E. Türeci, Physical reservoir computing using finitely-sampled quantum systems, [arXiv:2110.13849](https://arxiv.org/abs/2110.13849).
- [22] M. C. Mackey and L. Glass, Oscillation and chaos in physiological control systems, *Science* **197**, 287 (1977).
- [23] D. Stanev, N. Spagnolo, and F. Sciarrino, Deterministic optimal quantum cloning via a quantum-optical neural network, *Phys. Rev. Res.* **5**, 013139 (2023).
- [24] A. Motamedi, H. Zadeh-Haghighi, and C. Simon, Quantumness and learning performance in reservoir computing with a single oscillator, [arXiv:2304.03462](https://arxiv.org/abs/2304.03462).
- [25] G. R. Steinbrecher, J. P. Olson, D. Englund, and J. Carolan, Quantum optical neural networks, *Npj Quantum Inf.* **5**, 60 (2019).
- [26] H. Xu, T. Krisnanda, R. Bao, and T. C. H. Liew, The roles of Kerr nonlinearity in a bosonic quantum neural network, *New J. Phys.* **25**, 023028 (2023).
- [27] A. Bengtsson, P. Krantz, M. Simoen, I.-M. Svensson, B. Schneider, V. Shumeiko, P. Delsing, and J. Bylander, Nondegenerate parametric oscillations in a tunable superconducting resonator, *Phys. Rev. B* **97**, 144502 (2018).
- [28] C. Zhu, P. J. Ehlers, H. I. Nurdin, and D. Soh, Practical few-atom quantum reservoir computing, *Phys. Rev. Research* **7**, 023290 (2025).
- [29] M. Kornjača *et al.*, Large-scale quantum reservoir learning with an analog quantum computer, [arXiv:2407.02553](https://arxiv.org/abs/2407.02553).
- [30] D. Marković and J. Grollier, Quantum neuromorphic computing, *Appl. Phys. Lett.* **117**, 150501 (2020).
- [31] H.-Y. Huang, M. Broughton, J. Cotler, S. Chen, J. Li, M. Mohseni, H. Neven, R. Babbush, R. Kueng, J. Preskill, and J. R. McClean, Quantum advantage in learning from experiments, *Science* **376**, 1182 (2022).
- [32] Q. H. Tran and K. Nakajima, Learning temporal quantum tomography, *Phys. Rev. Lett.* **127**, 260401 (2021).
- [33] L. Balembois, *Magnetic resonance of a single electron spin and its magnetic environment by photon counting*, Theses, Université Paris-Saclay (2023).
- [34] The data and code that support this study are available in Zenodo with the identifier, <https://zenodo.org/records/15745370>.