

Physical reservoir computing using finitely-sampled quantum systems

Saeed A. Khan, Fangjun Hu, Gerasimos Angelatos, and Hakan E. Türeci

Department of Electrical and Computer Engineering, Princeton University, Princeton, NJ 08544

(Dated: November 5, 2021)

The paradigm of reservoir computing exploits the nonlinear dynamics of a physical reservoir to perform complex time-series processing tasks such as speech recognition and forecasting. Unlike other machine-learning approaches, reservoir computing relaxes the need for optimization of intra-network parameters, and is thus particularly attractive for near-term hardware-efficient quantum implementations. However, the complete description of practical quantum reservoir computers requires accounting for their placement in a quantum measurement chain, and its conditional evolution under measurement. Consequently, training and inference has to be performed using finite samples from obtained measurement records. Here we describe a framework for reservoir computing with nonlinear quantum reservoirs under continuous heterodyne measurement. Using an efficient truncated-cumulants representation of the complete measurement chain enables us to sample stochastic measurement trajectories from reservoirs of several coupled nonlinear bosonic modes under strong excitation. This description also offers a mathematical basis to directly compare the computational capabilities of a given physical reservoir operated across classical and quantum regimes. Applying this framework to the classification of quantum states of systems that are part of the same measurement chain as the quantum reservoir computer, we assess and explain measurement-contingent advantages and disadvantages of reservoir processing in quantum regimes. Our results also identify the vicinity of bifurcation points as presenting optimal nonlinear processing regimes of an oscillator-based quantum reservoir. The considered models are directly realizable in modern circuit QED experiments, while the framework is applicable to more general quantum nonlinear reservoirs.

I. INTRODUCTION

Reservoir computing is a machine-learning paradigm that emphasizes learning efficiency [1–4]: a linear combination of the accessible degrees of freedom of a physical ‘reservoir’ is the only quantity trained for a given task. Crucially, foregoing optimization of internal parameters does not necessarily hinder the computational capacity of reservoir computers, which in recent years have enabled efficient approaches to forecasting [5–7], inference [8, 9], control [10], and similar resource-intensive signal processing tasks [11–15]. By relaxing control over microscopic parameters of the learning machine, reservoir computing embraces the fundamental notion of computation with very general dynamical systems: physical or artificial dynamical systems stimulated by incoming time-dependent signals $\mathbf{u}(t)$ perform computation by transforming an input stream into an output stream, effectively computing a function $\mathcal{F}\{\mathbf{u}(t)\}$ [see Fig. 1(a)]. If the dynamical system is sufficiently complex and high-dimensional, the expressive power of even a linear classifier acting on the output state can enable approximation of a large set of functions $\mathcal{F}\{\cdot\}$ [16–18]. In this vein, diverse dynamical systems ranging from early mathematical black-box reservoirs [1] to physical systems across optical [19–24], electrical [25–28] and recently superconducting [29, 30] platforms have all been explored as reservoir computers [31–33].

This flexibility has naturally led to the examination of *quantum* systems as potential physical substrates for reservoir computing [34–39], with the promise of exponentially-scaling reservoir dimensionality. Early implementations of quantum reservoirs [see Fig. 1(a)] consider computation with explicitly quantum components (e.g. interacting spins, bosonic or fermionic degrees of freedom) and encode the input $\mathbf{u}(t)$ in the Liouvillian governing state dynamics, or even in the quantum reservoir state directly as a ‘quantum’ input. Reser-

voir output is often constructed assuming access to expectation values of reservoir dynamical variables that correspond loosely to measurable observables (e.g. occupation numbers via photodetectors, quadratures via homodyne measurement).

While such schemes have been used to theoretically show certain performance advantages compared to classical systems [40], critical foundational questions have yet to be addressed. The *experimentally-accessible* degrees of freedom of a quantum system are constrained given a specific input and output scheme, which must be taken into account. Importantly, the role of the quantum theory of measurement in sampling outputs from such a system for machine-learning remains unexplored. Furthermore, a quantum reservoir has yet to be treated from the perspective of a physically-motivated classical limit, the need for which is twofold. First, this provides the basis to unequivocally identify any possible advantages afforded by quantum reservoir processing. Secondly, the energy expended during a given computational task is a fundamental concern in physical computing; when considering computation with decreasing signal powers, the transition from classical to quantum regimes of reservoir operation must naturally be addressed.

This paper introduces a framework for quantum reservoir computing that allows the same physical reservoir to be controllably operated across quantum or classical regimes. To satisfy this requirement, we present an analysis built upon an intrinsically quantum-mechanical description of the complete measurement chain, including not only the physical reservoir but also inputs and measured outputs necessary to define a complete computational framework. The classical limit is then identified where the state of the measurement chain during computation is well-approximated as a product of coherent states. This unified approach allows us to define performance metrics (such as the *fidelity* of executing a particular task, or the *time-to-solution*) that account for limitations due

to a prescribed input-output scheme and measurement noise, and can be compared across classical or quantum regimes of operation of the same physical reservoir. Such a comparative metric provides the theoretical basis to identify task-specific and measurement-contingent advantages (or disadvantages, as we shall see) of operating physical reservoirs in the quantum regime.

The quantum-mechanical treatment of the complete measurement chain we require demands the specification of details often simplified in earlier studies. A representative layout of the quantum measurement chain we wish to describe is shown in Fig. 1(b), depicted here in the circuit QED (cQED) architecture [41, 42] (many components are generalizable to other platforms as well). The quantum reservoir computer (QRC) is modelled as a K -node quantum nonlinear device, whose driven-dissipative dynamics are governed by a Liouvillian superoperator \mathcal{L}_{QRC} . Inputs to the QRC must be transmitted via classical control fields used to interact with the quantum measurement chain; while classical inputs can be directly encoded into these control fields (shown here at a single frequency ω_d), the question of quantum inputs (also referred to as ‘quantum data’ in earlier work [36]) is more involved. We define quantum inputs as states of quantum systems (described by a Liouvillian superoperator \mathcal{L}_S) generated in the same cryogenic environment as the QRC. The flow of signals from this system to the reservoir is accurately modeled via a general coupling superoperator \mathcal{L}_c , which typically includes directional elements. This description of the composite measurement chain accurately accounts for correlations between the quantum system and the QRC, including the extreme case where the system and the quantum reservoir are *entangled*, a scenario with no analog in classical reservoir computing. Evolution of the QRC under classical or quantum inputs realizes a high-dimensional nonlinear mapping of the input to the reservoir’s state space.

Information encoded in the QRC state is extracted from the measurement chain via continuous heterodyne monitoring of the QRC nodes, which requires us to address limitations on measured outputs imposed by quantum theory. These measurements do not directly yield ensemble-averaged expectation values of reservoir state variables, but noisy output field measurement records that include quantum fluctuations nonlinearly transformed by the quantum reservoir itself, and are dependent upon the conditional quantum state of the measurement chain. All this is accounted for by formulating and solving for the conditional evolution of all dynamical elements in the quantum measurement chain under the prescribed measurement scheme (often called quantum trajectories [43] subject to measurement backaction). The final step in our quantum reservoir computing framework is an output layer applied to these measured outputs, performing *linear* operations that are trained on *finite* measurement records and not directly on expectation values. The latter in principle assumes access to an infinite number of measurement records, ignoring a key resource cost of quantum reservoir computing, and for higher-order expectation values can introduce *classical* nonlinear operations, whose computational influence should be distinguished from that of the nonlinear quantum reservoir

alone. This output layer can be implemented in software or, for real-time processing, via a customized closed-loop signal processing hardware such as FPGAs at room temperature.

Starting from the formal mathematical description of the complete measurement chain via a *stochastic master equation* (SME) [44, 45], we develop an efficient computational framework to simulate its use for quantum reservoir computing. By considering a representative machine-learning task of quantum state classification with small-scale QRCs, we demonstrate computation with quantum reservoirs on quantum inputs across classical and quantum regimes of reservoir operation. Our principal result is that advantages of operating reservoirs in the quantum regime are task-dependent: once limitations due to measurement are taken into account, not all state classification tasks immediately benefit from quantum reservoir processing. Specific tasks, which we identify, can be carried out with reduced measurement resources in quantum regimes of reservoir operation.

Before presenting the complete analysis of this quantum reservoir computing framework that constitutes the main body of this paper, in the remaining subsections of this introduction we preview the technical aspects that make this analysis possible, and a summary of its main results.

A. Technical approach

The evolution of the conditional quantum state $\hat{\rho}^c$ of the composite measurement chain under continuous measurement shown in Fig. 1(b) is formally described by the SME

$$d\hat{\rho}^c = \mathcal{L}_S \hat{\rho}^c dt + \mathcal{L}_c \hat{\rho}^c dt + \mathcal{L}_{\text{QRC}} \hat{\rho}^c dt + \mathcal{S}_{\text{meas}}(dW) \hat{\rho}^c \quad (1)$$

where $\mathcal{S}_{\text{meas}}(dW)$ is the stochastic measurement superoperator. For extended multi-mode and generally nonlinear networks of interest, such as the one presented by the QRC setup of interest in Fig. 1(b), Eq. (1) is typically unfeasible to simulate, limited by growing Hilbert space constraints. To accurately yet efficiently account for dynamics of the full measurement chain, a truncated-cumulants representation of its conditional quantum state is developed; the choice of representation is strongly dependent on the measurement scheme and Liouvillian dynamics in Eq. (1). This description extends earlier work on continuously-monitored linear modes [46, 47] to nonlinear multimode quantum systems, yielding a set of stochastic equations for measured observables that scales much more favourably with increasing system complexity. The accuracy of this representation requires specific conditions on the nonlinearity of the quantum reservoir and its operating conditions which, as discussed below, can be readily satisfied within the cQED architecture. Additionally, this framework admits a well-defined classical limit, where the joint quantum state of the measurement chain is well-approximated as a product of coherent states. Defined in this way, the classical limit can be analyzed without having to redefine the quantum-classical boundary (colloquially, the *Heisenberg-von-Neumann cut*) separating the measured system from the classical observer,

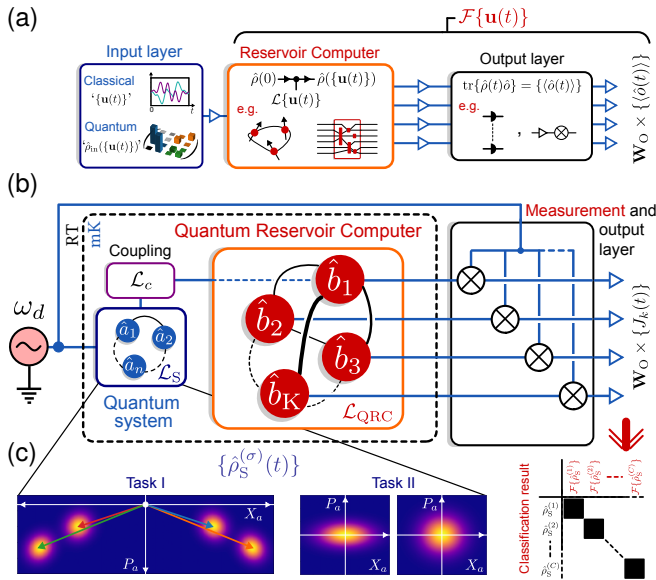


Figure 1. (a) General quantum reservoir computing frameworks. (b) Schematic representation of quantum reservoir computing framework under continuous measurement, describing a K -node quantum reservoir computer (QRC) as part of a measurement chain for processing signals encoded in classical control fields such as the single coherent drive (frequency ω_d), or generated by quantum systems in the same quantum environment. The dashed boundary separates classical controls at room temperature (RT) from quantum devices at cryogenic temperatures (mK). The quantum system is governed by Liouvillian \mathcal{L}_S , and is coupled to the QRC via \mathcal{L}_c . The QRC nodes undergo heterodyne measurement at the completion of the measurement chain, leading to its conditional evolution under measurement superoperators $\mathcal{S}_{\text{meas}}(dW)$. (c) The QRC is used to classify qubit pointer states generated in a cavity (task I), and states with equal phase and mean amplitude, but distinct variances (task II).

which we always place at the termination of the measurement chain.

B. Machine-learning tasks using quantum inputs

As a quantum information processing task, we consider the problem of classifying C states of a quantum system (governed by \mathcal{L}_S) embedded in the same measurement chain as the QRC. These states $\hat{\rho}_S^{(\sigma)}(t)$ (indexed by $\sigma = 1, \dots, C$) are dynamically generated by inputs to the measurement chain, and also evolve under the backaction of quantum measurement of the QRC nodes. As such, they are defined via the *conditional* quantum state of the measurement chain by tracing out the QRC, $\hat{\rho}_S^{(\sigma)}(t) = \text{tr}_{\text{QRC}}\{\hat{\rho}^{(\sigma)}(t)\}$. Note that this is simply a formal definition: our computational approach does not trace out any degrees of freedom of the measurement chain. For concreteness, we consider the classification of system states that (I) differ in their mean amplitudes and phases, and (II) differ only in their variances [see Fig. 1(c)]. Task I is closely related to the classification of pointer states for dispersive qubit readout in cQED [48], while Task II can be applied to distin-

guishing entangled and non-entangled quantum states [49]. In both cases, suitable physical QRCs can process these quantum inputs to produce distinct measured outputs, whose distributions include non-linearly transformed quantum fluctuations due to the action of the nonlinear QRC itself. To successfully implement the function $\mathcal{F}\{\hat{\rho}_S^{(\sigma)}(t)\} = \sigma$ for quantum state classification, as our framework does, the linear output layer must be trained to take these non-classical fluctuations into account.

C. Performance across classical and quantum operating regimes

The quantum reservoir computing framework presented here is constructed to assess computation fidelity as a function of quantum reservoir size K , and more importantly to explore the role of quantum resources such as entanglement [50] amongst the reservoir nodes. However, the generation and preservation of entanglement itself requires additional resources and operating constraints [51] on the QRC. As a first study, we find that comparing QRC performance across the transition from a classical regime of QRC operation to a quantum regime without necessarily demanding entanglement within the quantum reservoir already yields important insights into the role of quantum reservoir processing. Furthermore, this analysis provides the basis for future work to precisely identify any further advantages of reservoir entanglement.

We show that the input power to the measurement chain provides a natural, physically-motivated parameter to explore this classical-to-quantum transition: with decreasing input power, maintaining computational accuracy for quantum state classification requires increasing the QRC nonlinearity, which we show corresponds to the quantum regime of QRC operation. Importantly, across this transition, QRC performance on Tasks I and II shows important qualitative differences. Task I, which requires the QRC to distinguish the amplitudes and phases of quantum states, fares better in the classical regime. Task II, on the other hand, requires the QRC to distinguish information encoded in quantum fluctuations. We find that this task can be performed with reduced measurement resources in increasingly quantum regimes (for details see Sec. IV D, V D respectively).

D. Nonlinear processing near classical bifurcation points

While the two quantum state classification tasks probe distinct nonlinear processing capabilities of the QRC, we find that useful regimes of reservoir operation are unified by their proximity to classical bifurcation points. Near these bifurcation phases, we find an enhancement of the nonlinear response of the QRC to its inputs; this suggests a general signature to identify parameter regimes of optimal reservoir performance for a broader class of computational tasks.

E. Organization of the paper

The remainder of this paper is organized as follows. In Sec. II, we introduce the specific cQED realization of the K -node QRC and the measurement chain. In Sec. III, we use a minimal version of this measurement chain, a QRC comprised of a single nonlinear mode, to develop our theoretical approach to the simulation of its conditional dynamics, and discuss its regimes of validity. Sec. IV then considers the use of the QRC for classifying pointer states, and Sec. V for the classification of general Gaussian states with distinct variances (Tasks I and II respectively). In both cases, we provide a characterization of the distinct quantum measurement chains and their measured outputs, a detailed study of QRC parameters that enable successful classification, and a comparison of QRC performance across classical and quantum operating regimes.

II. PHYSICAL MODEL FOR A QUANTUM RESERVOIR COMPUTER

To describe the QRC, we adopt a model that is realizable within the cQED architecture, based on K coupled Kerr nonlinear modes furnished by Josephson junctions [52]. The nonlinear modes have frequencies $\{\omega_k\}$, nonlinearity strengths $\{\Lambda_k\}$, with linear coupling $\{g_{ij}\}$ between modes i and j , yielding the Liouvillian governing QRC dynamics,

$$\begin{aligned} \mathcal{L}_{\text{QRC}}\hat{\rho}^c = & -i\left[\sum_k \omega_k \hat{b}_k^\dagger \hat{b}_k - \sum_k \frac{\Lambda_k}{2} \hat{b}_k^\dagger \hat{b}_k^\dagger \hat{b}_k \hat{b}_k, \hat{\rho}^c\right] \\ & -i\left[\sum_{ij} g_{ij} (\hat{b}_i \hat{b}_j^\dagger + \hat{b}_i^\dagger \hat{b}_j), \hat{\rho}^c\right] + \sum_k \gamma_k \mathcal{D}[\hat{b}_k] \hat{\rho}^c. \end{aligned} \quad (2)$$

The linear damping is described by the standard dissipative superoperator $\mathcal{D}[\hat{o}] = \hat{o}\hat{\rho}\hat{o}^\dagger - \frac{1}{2}\{\hat{o}^\dagger\hat{o}, \hat{\rho}\}$, with rates $\{\gamma_k\}$. The decay channels for each node are continuously monitored with unit efficiency via a heterodyne measurement scheme, described by the stochastic measurement superoperator

$$\begin{aligned} \mathcal{S}_{\text{meas}}(dW)\hat{\rho}^c = & \sum_{k=1}^K \sqrt{\frac{\gamma_k}{2}} \left(\hat{b}_k \hat{\rho}^c + \hat{\rho}^c \hat{b}_k^\dagger - \langle \hat{b}_k + \hat{b}_k^\dagger \rangle^c \right) dW_k^X(t) + \\ & \sum_{k=1}^K \sqrt{\frac{\gamma_k}{2}} \left(-i\hat{b}_k \hat{\rho}^c + i\hat{\rho}^c \hat{b}_k^\dagger - \langle -i\hat{b}_k + i\hat{b}_k^\dagger \rangle^c \right) dW_k^P(t), \end{aligned} \quad (3)$$

where $dW_k^{X,P}$ are independent Wiener increments, and $\langle \hat{o} \rangle^c$ indicates the expectation value of an arbitrary operator \hat{o} with respect to the conditional quantum state, $\langle \hat{o} \rangle^c = \text{tr}\{\hat{\rho}^c \hat{o}\}$. The results of this continuous measurement are heterodyne cur-

rents $J_k^{X,P}(t)$ read out from each reservoir node,

$$J_k^X(t) dt = \sqrt{\frac{\gamma_k}{2}} \langle \hat{b}_k + \hat{b}_k^\dagger \rangle^c dt + dW_k^X(t), \quad (4)$$

$$J_k^P(t) dt = \sqrt{\frac{\gamma_k}{2}} \langle -i\hat{b}_k + i\hat{b}_k^\dagger \rangle^c dt + dW_k^P(t). \quad (5)$$

As described by Eq. (1), the QRC is just one part of the measurement chain. The signal-generating quantum system is described by the superoperator \mathcal{L}_S , and will be task-dependent; we will consider specific realizations throughout this paper. The coupling superoperator \mathcal{L}_c completes the measurement chain by allowing the measured quantum system and the QRC to interact. It can once again take an arbitrary form, allowing for both coherent and dissipative, reciprocal and non-reciprocal, and even nonlinear couplings. Depending on the type of coupling interaction, \mathcal{L}_c may also introduce local terms affecting the measured quantum system and the QRC; we will make these clear when they arise.

Having defined every element of the measurement chain in Fig. 1(b), we will now analyze the nature of its quantum dynamics, introducing an approach to studying this dynamics that is scalable to extended, multimode quantum systems.

III. QUANTUM DYNAMICS OF THE SINGLE-NODE KERR QRC

To analyze the nonlinear processing capabilities of this QRC, it will prove useful to begin with a simple minimal realization - a single Kerr node - and a single coherent drive incident on the QRC directly, via a separate input port with loss rate Γ . This situation describes QRC processing of classical signals encoded in the generally time-dependent coherent drive amplitude and phase. The resulting system schematic is shown in Fig. 2. We begin by analyzing the *unconditional* dynamics of the resulting system, described by the quantum master equation

$$d\hat{\rho} \equiv \mathcal{L}\hat{\rho} dt = \mathcal{L}_{\text{QRC}}\hat{\rho} dt - i[\hat{\mathcal{H}}_d, \hat{\rho}] dt \quad (6)$$

where the drive Hamiltonian in the frame rotating at the drive frequency ω_d is then

$$\hat{\mathcal{H}}_d = \eta(\hat{b} + \hat{b}^\dagger). \quad (7)$$

The Liouvillian \mathcal{L}_{QRC} describing the single-node QRC takes the form

$$\mathcal{L}_{\text{QRC}}\hat{\rho} = -i\left[-\Delta\hat{b}^\dagger\hat{b} - \frac{\Lambda}{2}\hat{b}^\dagger\hat{b}^\dagger\hat{b}\hat{b}, \hat{\rho}\right] + (\gamma + \Gamma)\mathcal{D}[\hat{b}]\hat{\rho} \quad (8)$$

where $\Delta = \omega_d - \omega$ is the detuning of the drive frequency from the bare QRC node frequency ω . For simplicity of notation, in what follows we introduce for convenience the total linear damping $\tilde{\gamma} = \gamma + \Gamma$ of the nonlinear mode.

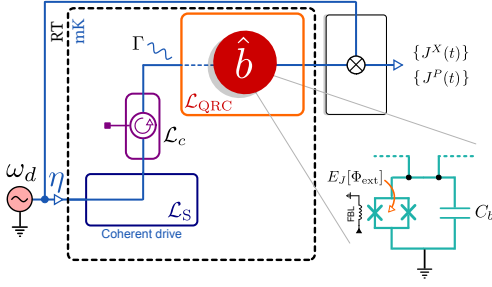


Figure 2. Schematic showing the fundamental node of the proposed QRC as part of a simplified measurement chain: a ‘system’ described by a single frequency coherent drive incident on the QRC node via an input channel that is not measured, and continuous weak measurement of the QRC node output. The QRC node can be realized in cQED as a tunable-frequency Kerr oscillator using a capacitively shunted SQUID loop [53].

A. Equations of motion approach based on truncation of higher-order cumulants

The quantum state described by an arbitrary density matrix $\hat{\rho}$ can always be written in terms of normal-ordered cumulants, defined for a single quantum mode \hat{b} as

$$C_{b^\dagger p b^q} \equiv \frac{\partial^{p+q}}{\partial (iz^*)^p \partial (iz)^q} \ln \text{tr} \{ \hat{\rho} e^{iz^* \hat{b}^\dagger} e^{iz \hat{b}} \} \Big|_{z=z^*=0}. \quad (9)$$

The term in curly braces may immediately be recognized as the characteristic function of $\hat{\rho}$ [54], associated with normal-ordered phase-space distributions such as the P -representation; this connection is elaborated upon in Appendix A. The above definition generalizes straightforwardly to states $\hat{\rho}$ describing multiple bosonic modes.

Cumulants are defined by their order $p + q$, and provide an equivalent set of variables to describe the dynamics of a quantum state. Specifically, we can write down the equation of motion for the first-order cumulant of the single-node QRC ($C_b \equiv \langle \hat{b} \rangle = \text{tr} \{ \mathcal{L} \hat{\rho} \hat{b} \}$),

$$\begin{aligned} \dot{\langle \hat{b} \rangle} &= \left(i\Delta - \frac{\tilde{\gamma}}{2} \right) \langle \hat{b} \rangle + i\Lambda \langle \hat{b}^\dagger \rangle \langle \hat{b} \rangle \langle \hat{b} \rangle - i\eta \\ &+ i\Lambda \left(C_{bb} \langle \hat{b}^\dagger \rangle + 2C_{b^\dagger b} \langle \hat{b} \rangle + C_{b^\dagger bb} \right) \end{aligned} \quad (10)$$

where $C_{bb} = \langle \hat{b} \hat{b} \rangle - \langle \hat{b} \rangle^2$, $C_{b^\dagger b} = \langle \hat{b}^\dagger \hat{b} \rangle - \langle \hat{b}^\dagger \rangle \langle \hat{b} \rangle$, and $C_{b^\dagger bb} = \langle \hat{b}^\dagger \hat{b} \hat{b} \rangle - 2\langle \hat{b}^\dagger \hat{b} \rangle \langle \hat{b} \rangle - \langle \hat{b} \hat{b} \rangle \langle \hat{b}^\dagger \rangle + 2\langle \hat{b}^\dagger \rangle \langle \hat{b} \rangle^2$. Expressions for higher-order cumulants become increasingly more unwieldy, but can be systematically obtained, as discussed in Appendix C. We see immediately that the first-order cumulant $\langle \hat{b} \rangle$ couples to cumulants of second and higher order. We can similarly obtain equations of motion for unique second-order

cumulants $C_{b^\dagger b}, C_{bb}$,

$$\begin{aligned} \dot{C}_{b^\dagger b} &= -\tilde{\gamma} C_{b^\dagger b} - i\Lambda (C_{bb} \langle \hat{b}^\dagger \rangle^2 - C_{bb}^* \langle \hat{b} \rangle^2) \\ &- i\Lambda (C_{b^\dagger bb} \langle \hat{b} \rangle - C_{b^\dagger bb}^* \langle \hat{b}^\dagger \rangle), \end{aligned} \quad (11a)$$

$$\begin{aligned} \dot{C}_{bb} &= \left(i2\Delta - \tilde{\gamma} + i\Lambda (1 + 4|\langle \hat{b} \rangle|^2 + 6C_{b^\dagger b}) \right) C_{bb} \\ &+ i\Lambda \langle \hat{b} \rangle^2 (1 + 2C_{b^\dagger b}) + i2\Lambda (C_{b^\dagger bbb} + 2\langle \hat{b} \rangle C_{b^\dagger bb} + C_{bbb} \langle \hat{b}^\dagger \rangle). \end{aligned} \quad (11b)$$

Clearly, second-order cumulants couple to cumulants of third and fourth order. This is a general feature of nonlinear quantum systems: moments and thus cumulants of a certain order can couple to those of higher-order, leading to an infinite hierarchy of equations that do not form a closed set.

However, an infinite set of normal-ordered cumulants is not necessarily required to describe all multimode quantum states. In particular, a multimode quantum system in a product of coherent states is described entirely by its nonzero first-order cumulants; all cumulants of order $p + q > 1$ vanish (see Appendix B for derivations). Multimode quantum states that are defined entirely by their first and second-order cumulants admit Gaussian phase-space representations, and are labelled Gaussian states. States with nonzero cumulants of third or higher-order are thus by definition non-Gaussian states [55, 56].

Our numerical approach leverages this efficient representation of specific multimode quantum states in terms of cumulants. In particular, we consider an ansatz wherein the quantum state of the complete measurement is described entirely by cumulants up to a finite order $p + q \leq n_{\text{trunc}}$; all cumulants of order $p + q > n_{\text{trunc}}$ are thus set to zero, truncating the hierarchy and yielding a closed set of equations for the retained nonzero cumulants. In this paper, we choose $n_{\text{trunc}} = 2$ for a quantum measurement chain defined entirely by its first and second-order cumulants, although the truncation can similarly be carried out at higher order. The resulting Truncated Equations Of Motion (TEOMs), and their generalization to *conditional* dynamics of multimode quantum systems introduced in Sec. III C, form the basis of our computational approach in this paper.

Since we are specifically interested in *nonlinear* quantum systems for reservoir computing, which can generate higher-order cumulants in dynamics, one must ask when such an ansatz may hold. In the two following subsections, we identify system-specific drive and parameter regimes where this ansatz is valid, and verify our conclusions via numerical comparisons with exact results and (S)ME simulations.

B. Application to unconditional QRC node dynamics

We will use the single coherently-driven QRC node modeled as a Kerr oscillator to benchmark the TEOMs approach. In doing so, we make extensive use of the exact quantum steady-state solution of the coherently-driven Kerr oscillator, obtained via the complex- P representation [57]. This provides access to steady-state cumulants of arbitrary order (see Appendix E for details), which can be used to identify regions

of parameter space where a truncation of cumulants beyond second-order can be justified.

One such regime should be furnished by the *classical limit* of the single Kerr QRC node, where the QRC is always in a coherent state and thus described entirely by its first-order cumulants. To identify this classical limit, and thus explore deviations from it due to quantum effects, it proves convenient to work with specific scaled quantities: after introducing dimensionless time $t' = \tilde{\gamma}t$ and energy scales $(\Delta', \Lambda', \eta') = (\Delta, \Lambda, \eta)/\tilde{\gamma}$, we scale $\langle \hat{b} \rangle' \rightarrow \sqrt{\Lambda'} \langle \hat{b} \rangle$, and impose our truncated ansatz, following which Eq. (10) becomes

$$\begin{aligned} \frac{d}{dt'} \langle \hat{b} \rangle' &= \left(i\Delta' - \frac{1}{2} \right) \langle \hat{b} \rangle' + i \langle \hat{b}^\dagger \rangle' \langle \hat{b} \rangle' \langle \hat{b} \rangle' - i\mathcal{N} \\ &+ i\Lambda' \left(C_{bb} \langle \hat{b}^\dagger \rangle' + 2C_{b^\dagger b} \langle \hat{b} \rangle' \right). \end{aligned} \quad (12)$$

The first line in Eq. (12) makes no reference to second-order cumulants of the QRC node state; it thus describes dynamics when the QRC is in a coherent state (for which such cumulants vanish), corresponding to the classical limit. Dynamics in this limit depend on the drive and the nonlinearity via a single dimensionless parameter [58, 59],

$$\mathcal{N} = \eta' \sqrt{\Lambda'} = \frac{\eta}{\tilde{\gamma}} \sqrt{\frac{\Lambda}{\tilde{\gamma}}}. \quad (13)$$

The second line of Eq. (12) does however depend on cumulants; these terms describe quantum dynamics that lead to a deviation from purely coherent state evolution. In these scaled units, the TEOMs for the second-order cumulants obtained from Eq. (11a), (11b) take the form

$$\dot{C}_{b^\dagger b} = -C_{b^\dagger b} - i(C_{bb} \langle \hat{b}^\dagger \rangle'^2 - C_{bb}^* \langle \hat{b} \rangle'^2), \quad (14a)$$

$$\begin{aligned} \dot{C}_{bb} &= \left(i2\Delta' - 1 + i(\Lambda' + 4|\langle \hat{b} \rangle'|^2 + 6\Lambda' C_{b^\dagger b}) \right) C_{bb} \\ &+ i \langle \hat{b} \rangle'^2 (1 + 2C_{b^\dagger b}). \end{aligned} \quad (14b)$$

Consider now the transformation where $\Lambda' \rightarrow 0$, while \mathcal{N} is held fixed (by increasing the drive strength η' simultaneously). Terms in the first line of Eq. (12) describing classical dynamics remain unchanged, while those in the second line describing quantum deviations become smaller, *provided* second-order cumulants do not grow with Λ' . From Eqs. (14a), (14b), it is clear that $C_{b^\dagger b}$ is independent of Λ' , while C_{bb} has a dependence that is negligible as $\Lambda' \rightarrow 0$ for $|\langle \hat{b} \rangle'| \neq 0$. Hence by keeping \mathcal{N} fixed while decreasing the nonlinearity strength, the influence of second-order cumulants on dynamics of the QRC node amplitude is suppressed, describing the classical limit of QRC operation. Physically, this transformation leads to an increase in the unscaled drive η' and QRC node amplitude $|\langle \hat{b} \rangle|$, and hence occupation number, decreasing the relative impact of quantum fluctuations in agreement with conventional notions of classicality. Conversely, increasing nonlinearity strength for fixed \mathcal{N} should enhance the impact of quantum fluctuations, marked by a systematic increase in second-order cumulants, and taking the Kerr QRC gradually towards a quantum regime of operation.

To explore this transition, we start with the classical limit of the QRC as determined by terms in the first line of Eq. (12), for which the steady-state phase diagram can be easily computed (see Appendix G). This phase diagram, shown in (Δ, \mathcal{N}) -space in the center panel of Fig. 3(a), is the well-known result for a coherently-driven classical Kerr oscillator [58, 59]: the orange region depicts the classical bistability, which emerges for sufficiently negative detuning given our choice of the sign of the Kerr nonlinearity.

This bistability is a feature specific to nonlinear systems; in its vicinity, we can expect an enhanced nonlinear response of the QRC observables to QRC inputs. This is explicitly observed in the scaled field amplitude $|\langle \hat{b} \rangle'|$, plotted in Fig. 3(b) as a function of \mathcal{N} for fixed $\Delta/\gamma = -1.0$ (that is, across the indicated vertical dashed line of the classical phase diagram). The solid blue curve is the exact complex- P result for $\Lambda/\gamma = 0.005$, while blue dots are obtained by evolving the TEOMs, Eqs. (10), (11a), (11b), to their steady state numerically. For small \mathcal{N} , the response follows that of a linear oscillator (dashed black line), but becomes nonlinear with increasing \mathcal{N} . Also shown in dashed orange is the response of the classical nonlinear oscillator (explicitly dropping second-order cumulants), which does not depend on the actual nonlinearity strength but only on \mathcal{N} .

The influence of nonlinearity is similarly evident when plotting $|\langle \hat{b} \rangle'|$ as a function of Δ , for fixed $\mathcal{N} = 0.385$ in Fig. 3(c) (across the horizontal dashed line of the classical phase diagram). The black dashed curve indicates the standard Lorentzian response of a linear oscillator, centered around the bare frequency ω_b . The response of the nonlinear QRC node is clearly that of a Lorentzian deformed towards lower frequencies due to the Kerr-induced frequency shift.

In both results, we find excellent agreement between the exact quantum results, the TEOMs, *and* classical nonlinear dynamics. This is indicative of operating parameters where not only is a truncated description using first and second-order cumulants sufficient, but the influence of second-order cumulants on first-order cumulants is negligible as well. Consequently, we expect the QRC node here to function as a classical reservoir node, still allowing for nonlinear processing of information encoded in the drive amplitude, phase, and/or frequency.

If we now keep \mathcal{N} fixed but increase the nonlinearity strength Λ , the QRC can be moved to a quantum regime of operation. We plot the QRC response as a function of \mathcal{N} for $\Delta/\gamma = -1.0$ in Fig. 3(b), and as a function of Δ for $\mathcal{N} = 0.385$ in Fig. 3(c), for three different nonlinearity strengths $\Lambda/\gamma \in [0.005, 0.02, 0.05]$. In both cases, we see that with increasing nonlinearity, the response deviates from that of the classical nonlinear oscillator (which recall is essentially equivalent to the quantum result for weak $\Lambda/\gamma = 0.005$). From Eq. (12), it is clear that this deviation is due to the coupling of first-order moments to second-order cumulants; this is precisely the effect accounted for by the TEOMs, which thus agree with the full quantum result.

Also shown in Figs. 3(d) and (e) are steady-state cumulants plotted across the same cross-sections as Figs. 3(b) and (c) respectively, with panels from left to right indicating

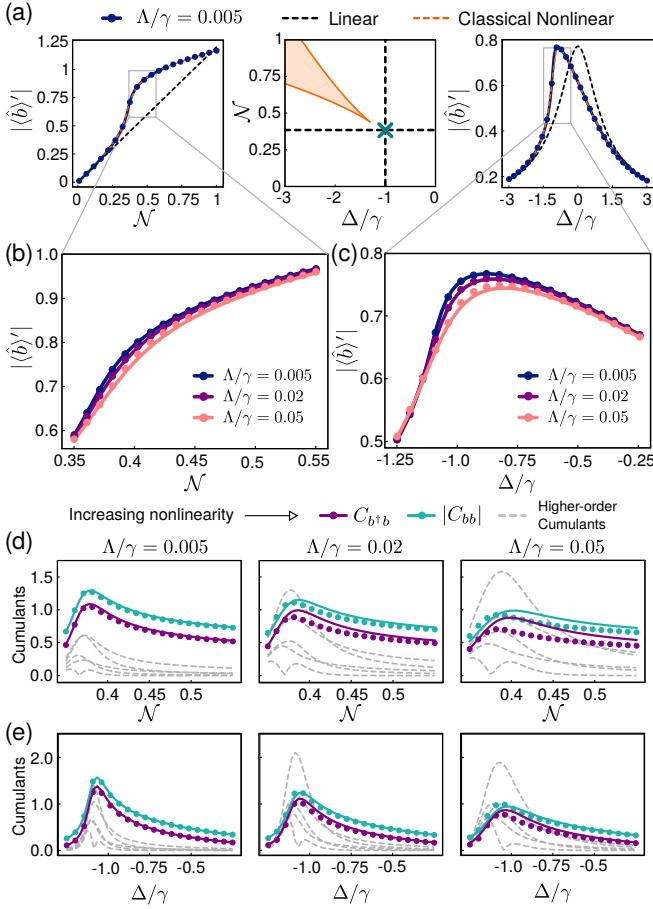


Figure 3. Classical and quantum dynamics of a single Kerr QRC node. In this plot, colored solid curves are exact complex- P results, colored dots are TEOMs simulations in the long-time limit, and dashed orange curves are classical results. Other curves are specified when relevant. (a) Center panel: *Classical* phase diagram of the single coherently-driven Kerr QRC node. Left panel: QRC amplitude as a function of \mathcal{N} for fixed detuning $\Delta/\gamma = -1.0$ (across indicated vertical cross-section of the phase diagram). Right panel: QRC amplitude as a function of detuning Δ for fixed $\mathcal{N} = 0.385$ (across indicated horizontal cross-section of the phase diagram). Dashed black lines in both plots indicate the response for a linear oscillator. (b) Steady-state scaled amplitude $|\langle \hat{b} \rangle|$ for nonlinearity strengths $\Lambda/\gamma \in [0.005, 0.02, 0.05]$, as a function of \mathcal{N} for fixed $\Delta/\gamma = -1.0$ (region within the gray box in (a), left panel). (c) Same as (b) but as a function of detuning Δ for fixed $\mathcal{N} = 0.385$ (region within the gray box in (a), right panel). (d) Steady-state second-order cumulants ($C_{b^\dagger b}$, purple, and $|C_{bb}|$, green) and higher (third, fourth) order cumulants (dashed gray) across the cross-section in (b); nonlinearity increases from left to right. (e) Same as (d) but now for the cross-section in (c).

stronger nonlinearities. Second-order cumulants are plotted in green and purple (lines are exact complex- P results, dots are TEOMs); they clearly display non-monotonic behaviour as a function of \mathcal{N} and Δ , acquiring large magnitudes for specific operating parameters in the vicinity of the classical bistability. Physically, second-order cumulants are related to the maximum and minimum (dimensionless) quadrature variances of

the QRC nodes: $\frac{1}{2} + C_{b^\dagger b} \pm |C_{bb}|$ respectively, and their magnitudes are thus measures of amplification and squeezing of quantum fluctuations due to the underlying Kerr nonlinearity.

Third and fourth order cumulants are shown in the same plots in dashed gray (exact complex- P only). There are two features of note. Firstly, the overall magnitude of these higher-order cumulants relative to second-order cumulants appears to increase for stronger nonlinearities. This is an indicator of the emergence of non-Gaussian features with increasing Λ , as expected, and generally leads to greater deviation between TEOMs and exact quantum results for first and second-order cumulants. Secondly, even for a fixed nonlinearity strength, the magnitude of cumulants depends on operating parameters (\mathcal{N} and Δ) that lead to operation in the vicinity of the classical bistability. Hence, while results obtained from TEOMs for such operating parameters can deviate from the exact quantum solution, away from these regions the agreement improves, even for strong nonlinearities.

C. Conditional QRC node dynamics

We have so far used steady-state, unconditional quantities to analyze the nonlinear processing capabilities of a single Kerr-based QRC node, and identified parameter regimes where an approach based on truncated EOMs is valid. However, to allow for processing of time-dependent information by the QRC, and to analyze the output from the QRC obtained as *individual* measurement records, an approach is needed that is able to accurately capture the *conditional, dynamical* evolution of the QRC node. This is described by the SME

$$d\hat{\rho}^c = \mathcal{L}\hat{\rho}^c dt + \mathcal{S}_{\text{meas}}(dW)\hat{\rho}^c \quad (15)$$

where $\mathcal{S}_{\text{meas}}(dW)$ is defined in Eq. (3), and our current system assumes the $K = 1$ case. Expectation values with respect to the conditional quantum state under measurement $\hat{\rho}^c$ are thus also conditional, which we indicate by the superscript c . For example, the equation of motion for $\langle \hat{b} \rangle^c$ is now given by $\text{tr}\{(\mathcal{L} dt + \mathcal{S}_{\text{meas}}(dW))\hat{\rho}^c \hat{b}\}$, which takes the form (see Appendix D for details)

$$d\langle \hat{b} \rangle^c = \text{tr}\{(\mathcal{L}\hat{\rho}^c \hat{b})\} dt + \sqrt{\frac{\gamma}{2}} (C_{b^\dagger b}^c + C_{bb}^c) dW^X(t) + i\sqrt{\frac{\gamma}{2}} (C_{b^\dagger b}^c - C_{bb}^c) dW^P(t). \quad (16)$$

The first line $\propto \mathcal{L}$ simply includes terms from Eq. (10), with expectation values replaced by their conditional counterparts. Terms in the second line are due to the measurement superoperator, and render the equation of motion for $\langle \hat{b} \rangle^c$ *stochastic*, conditioned on the trajectory-specific realizations of $dW^{X,P}(t)$. Dropping cumulants higher than second order in Eq. (16), arising in the first line, yields the *stochastic* truncated EOMs (STEOMs) for $\langle \hat{b} \rangle^c$.

We can similarly obtain truncated equations of motion for the conditional cumulants $C_{bb}^c, C_{b^\dagger b}^c$ under heterodyne mea-

surement:

$$dC_{b^\dagger b}^c = \left[\text{tr}\{\mathcal{L}\hat{\rho}^c\hat{b}^\dagger\hat{b}\} - \langle\hat{b}^\dagger\rangle^c \text{tr}\{\mathcal{L}\hat{\rho}^c\hat{b}\} - \langle\hat{b}\rangle^c \text{tr}\{\mathcal{L}\hat{\rho}^c\hat{b}^\dagger\} \right] dt - \gamma \left[(C_{b^\dagger b}^c)^2 + C_{bb}^c (C_{bb}^c)^* \right] dt, \quad (17a)$$

$$dC_{bb}^c = \left[\text{tr}\{\mathcal{L}\hat{\rho}^c\hat{b}\hat{b}\} - 2\langle\hat{b}\rangle^c \text{tr}\{\mathcal{L}\hat{\rho}^c\hat{b}\} \right] dt - 2\gamma C_{bb}^c C_{b^\dagger b}^c dt. \quad (17b)$$

Again, terms $\propto \mathcal{L}$ are as found in Eqs. (11a), (11b) post-truncation, with expectation values replaced by their conditional counterparts. The second line of each equation describes the evolution due to measurement, which at first glance appears deterministic: note that Wiener increments $dW^{X,P}(t)$ make no appearance. For *linear* quantum systems under continuous weak measurement, this is in fact the case: second-order cumulants form a closed set described by the above equations, and no stochastic terms arise [46, 47, 60, 61]. However, for *nonlinear* quantum systems of interest here, second-order cumulants can couple to the *stochastic* first-order moments (here, via terms $\propto \mathcal{L}$), rendering the conditional evolution of second-order cumulants generally stochastic as well.

Eqs. (16), (17a), (17b) define the STEOMs for the single Kerr QRC node. To assess their validity, we compare their simulation results against integration of the SME, Eq. (15). For concreteness, we consider the operating point marked by the teal ‘X’ in the phase diagram of Fig. 3(a): $\mathcal{N} = 0.385$, $\Delta/\gamma = -1.0$. Note that the same noise realizations $dW^{X,P}(t)$ are used for both simulations to ensure the compared trajectories are conditioned on the same measurement record. Results are included in Fig. 4(a), (b) for $\Lambda/\gamma \in [0.005, 0.02]$ respectively, showing first-order moments (left panel) and second-order cumulants (center panel), with STEOMs results in white, and SME results obtained using QuTiP [62] in color. Both methods agree very well, especially for first-order cumulants. The increasing fluctuations in the second-order cumulants are a direct signature of the non-linearity of the quantum system under study.

However, in a real experiment, such information about the QRC state must be extracted from measurement records. For the case of heterodyne measurement, the obtained records $J^{X,P}(t)$ are defined in Eqs. (4), (5). These records are typically processed to reduce noise, most commonly via temporal filtering; the processed records define *measured quadratures*

$$I^{X,P}(t) = \frac{1}{t} \int_0^t d\tau J^{X,P}(\tau) \quad (18)$$

which here describe processing of single-shot readout records. Measured quadratures obtained using a single set of measurement records are shown in the *measured phase space* given by $(I^X(t), I^P(t))$ in the right panel of Fig. 4(a), (b) $\Lambda/\gamma \in [0.005, 0.02]$ as before; SME results are shown in solid black, with the STEOMs results plotted on top, using a colorscale going from black to yellow indicating increasing time. Both show excellent agreement; this is unsurprising since $J^{X,P}(t)$ yield information about the underlying conditional dynamics of $\langle\hat{b}\rangle^c$, which also agree very well between the two methods.

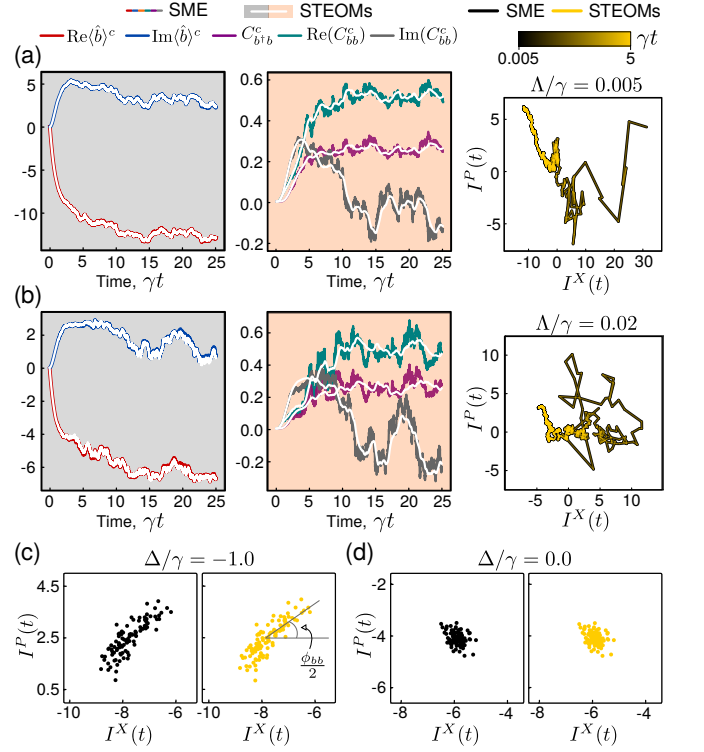


Figure 4. Benchmarking STEOMs against full SME simulations. Simulated parameters are as marked by the teal ‘X’ in the phase diagram of Fig. 3(a): $\mathcal{N} = 0.385$, $\Delta/\gamma = -1.0$, for (a) $\Lambda/\gamma = 0.005$ and (b) $\Lambda/\gamma = 0.02$. Left panel shows real and imaginary parts of the conditional expectation value $\langle\hat{b}\rangle^c$, and center panel shows conditional second-order cumulants $C_{bb}, C_{b^\dagger b}$; colored (white) curves are SME (STEOMs) simulations. Right panel shows measured quadratures $I^X(t), I^P(t)$: the SME result is plotted in solid black, while the STEOMs result is plotted in yellow, with a lighter shade indicating earlier times. (c) Measured quadratures $\{I^X(t), I^P(t)\}$ at $\gamma t = 20.0$ obtained using SME (left panel) and STEOMs (right panel), for $\mathcal{N} = 0.385$, $\Lambda/\gamma = 0.02$, and $\Delta/\gamma = -1.0$. (d) Same as (c) but for $\Delta/\gamma = 0.0$.

Note that the measured quadratures are themselves stochastic quantities, whose statistics are correlated with the underlying measured system state. This can be seen via the measured quadrature distributions in Fig. 4(c), obtained by simulating several measurement records (here 100 in total) and plotting $\{I^X(t), I^P(t)\}$ at $\gamma t = 20.0$, for $\Delta/\gamma = -1.0$ as before. At this operating point near the classical bistable region, the Kerr nonlinearity amplifies the magnitude of cumulants $C_{bb} \equiv |C_{bb}|e^{i\phi_{bb}}$ [see Fig. 3 (d)] leading to squeezing of the internal QRC field along the axis determined by $\phi_{bb}/2$ (and amplification along the orthogonal quadrature). This internal QRC squeezing manifests as squeezing of the measured quadrature distribution; the squeezing axis is unchanged since the temporal filter defining measured quadratures via Eq. (18) is quadrature-agnostic, and thus only reduces the overall noise power while preserving its relative strength amongst measured quadratures. On the other hand, for $\Delta/\gamma = 0.0$, further away from the bistable region, C_{bb} is much smaller in magnitude and the measured quadrature distribution in Fig. 4(d) exhibits

no squeezing. Excellent agreement between SME (left) and STEOMs (right) is observed for both cases.

D. Validity and scalability of (S)TEOMs

The key observations from our analysis of Kerr QRC operating regimes and the generation of higher-order cumulants (summarized in Fig. 3) are twofold. Primarily, the magnitude of nonzero higher-order cumulants increases with the strength of the Kerr nonlinearity relative to the QRC node damping rate, $\Lambda/\tilde{\gamma}$, which coincides with the QRC under coherent driving transitioning from classical to quantum regimes. However, operating regimes in the vicinity of the classical bistability also lead to an enhancement in higher-order cumulants, regardless of nonlinearity strength.

To be confident of the validity of the truncated cumulants approach, we therefore analyze QRCs with nonlinearity strengths $\Lambda/\tilde{\gamma} \leq 0.02$, and operate *near* but not within the classically-bistable region. This truncated cumulants approach provides a highly efficient mathematical description when applied to multimode quantum systems. For a measurement chain comprising an N -mode quantum system coupled to a K -node QRC, the composite system of $N_T = N + K$ total quantum modes is described by $2N_T^2 + 3N_T$ unknowns using (S)TEOMs, scaling quadratically with N_T instead of the exponential growth in required Hilbert space size for full (S)ME simulations. This description also places no constraints on modal occupation numbers, enabling our exploration of a well-defined classical limit of the coherently-driven Kerr QRC. Our (S)TEOMs approach is built to be a scalable theoretical framework via an efficient computer-algebra implementation, which allows the calculation of equations of motion of retained cumulants up to order n_{trunc} (here, second-order) for nonlinear quantum systems comprising arbitrary numbers of bosonic modes under continuous measurement. Supplementary benchmarking simulations for systems with more than one mode are included in Appendix F.

In the remainder of this paper, we employ (S)TEOMs to simulate dynamics of measurement chains comprising measured quantum systems coupled to multimode nonlinear QRCs. We find that such QRCs are able to successfully perform the quantum state classification tasks we have set out as objectives in the introduction, while operating within the aforementioned constraints.

IV. CAVITY MEASUREMENT: CLASSIFYING QUBIT POINTER STATES

We are now in a position to apply our quantum reservoir computing framework to the task of quantum state classification, beginning with Task I as identified in Fig. 1(c): classifying Gaussian states with mean values that differ in amplitude or phase. A useful example of this task is the classification of pointer states generated during cavity-mediated dispersive-qubit measurements [48, 63], when qubit evolution during measurement is negligible. Classification for nontrivial

qubit evolution during measurement has been considered using classical reservoirs in Ref. [30]. We begin with a description of all elements of the measurement chain included in Eq. (1) for this task, depicted schematically in Fig. 5(a).

A. Hardware measurement chain

1. Single cavity mode and nonreciprocal coupling to QRC

Our measured quantum system is described by the Liouvillian superoperator \mathcal{L}_S defined as

$$\mathcal{L}_S \hat{\rho}^c = -i[-\Delta_a^{(\sigma)} \hat{a}_1^\dagger \hat{a}_1 + \eta(\hat{a}_1 + \hat{a}_1^\dagger), \hat{\rho}^c] + \kappa \mathcal{D}[\hat{a}_1] \hat{\rho}^c \quad (19)$$

which defines a single cavity mode \hat{a}_1 , with frequency ω_{a_1} , damping rate κ , and coherently driven at frequency ω_d with drive strength η . The superoperator is written in the frame rotating at the drive frequency, rendering the drive term time-independent and introducing the cavity mode detunings

$$\Delta_a^{(\sigma)} = \omega_d - (\omega_{a_1} + \delta^{(\sigma)}). \quad (20)$$

Here frequency shifts $\delta^{(\sigma)}$ indexed by σ model the case of dispersive qubit readout: for the measurement of two qubits $\hat{\sigma}_{1,2}^z$ dispersively coupled to the cavity with strengths $\chi_{1,2}$ respectively, the dispersive shifts corresponding to the four joint qubit \hat{z} -basis eigenstates are given by $\delta^{(\sigma)} \in \{\chi_1 - \chi_2, \chi_1 + \chi_2, -\chi_1 - \chi_2, -\chi_1 + \chi_2\}$. Under coherent driving, the cavity state evolves to a distinct amplitude and phase for each dispersive shift, serving as a pointer to the corresponding qubit state. The computational task set for the QRC is then to distinguish these pointer states. For concreteness, we choose the drive to be resonant with the bare measured cavity frequency, $\omega_d = \omega_a$, and choose cavity detunings that map to effective dispersive shifts $\chi_1 = 2\kappa, \chi_2 = 0.5\kappa$.

The description of this quantum system alone is insufficient to define its conditional dynamics when embedded within a measurement chain, which includes the measured QRC and a coupling element. The specific form of the coupling is important as well: ideally, we require the information from the measured quantum system to be seen by the QRC, but for the QRC to not strongly couple to, and thus renormalize the properties of, the measured quantum system.

These requirements can be met by ensuring that \mathcal{L}_c describes a non-reciprocal coupling between the measured quantum system and the QRC. Here we couple the single cavity mode \hat{a}_1 to a single node \hat{b}_1 of the QRC via a directional amplifier interaction, realized by balancing a coherent QND interaction with its appropriately chosen dissipative counterpart [64]. This is defined by the specific coupling superoperator \mathcal{L}_c defined as

$$\mathcal{L}_c \hat{\rho} = -i[-g_c \hat{P}_1 \hat{X}_a, \hat{\rho}] + \Gamma_c \mathcal{D}[\hat{X}_a + i\hat{P}_1] \hat{\rho} \quad (21)$$

where $\hat{X}_a = \frac{1}{\sqrt{2}}(\hat{a}_1 + \hat{a}_1^\dagger)$, $\hat{P}_a = -\frac{i}{\sqrt{2}}(\hat{a}_1 - \hat{a}_1^\dagger)$ are the canonically conjugate quadratures of the measured cavity mode, and

\hat{X}_1, \hat{P}_1 are analogously-defined quadratures of the QRC \hat{b}_1 mode. The form of the interaction, realizable in cQED using multiple parametric pumps [65, 66], ensures that the quadratures \hat{P}_a, \hat{X}_1 undergo coupled evolution, as defined by the equations of motion

$$\frac{d}{dt}\langle\hat{P}_a\rangle^c = \text{tr}\{[\mathcal{L}_S + \mathcal{S}_{\text{meas}}(dW)]\hat{\rho}^c\hat{P}_a\} - [g_c - \Gamma_c]\langle\hat{P}_1\rangle^c, \quad (22a)$$

$$\frac{d}{dt}\langle\hat{X}_1\rangle^c = \text{tr}\{[\mathcal{L}_{\text{QRC}} + \mathcal{S}_{\text{meas}}(dW)]\hat{\rho}^c\hat{X}_1\} - [g_c + \Gamma_c]\langle\hat{X}_a\rangle^c \quad (22b)$$

The first term in each equation describes the uncoupled dynamics of the corresponding system, while the second term includes the conditional evolution due to the measurement superoperator. The third term is the coupling contribution: under the specific choice of coupling strength and phase

$$g_c = \Gamma_c \quad (23)$$

it is clear from Eq. (22b) that \mathcal{L}_c realizes a directional amplifier: the cavity \hat{X}_a quadrature is amplified by a factor $\propto \Gamma_c$ and impinges on the QRC, while the cavity itself is not driven by the QRC in the opposite direction.

Even if the cavity-QRC coupling is nonreciprocal, measurement of the QRC nodes downstream leads to measurement-conditioned evolution of the cavity $\propto \mathcal{S}_{\text{meas}}(dW)$ [67]. This is evident from the conditional equation of motion for the cavity $\langle\hat{X}_a\rangle^c$ quadrature (which drives the QRC),

$$\begin{aligned} \frac{d}{dt}\langle\hat{X}_a\rangle^c &= \text{tr}\{\mathcal{L}_S\hat{\rho}^c\hat{X}_a\} \\ &+ \sum_k \frac{\sqrt{\gamma_k}}{2} [C_{ab_k}^c + (C_{ab_k}^c)^* + C_{a^\dagger b_k}^c + (C_{a^\dagger b_k}^c)^*] dW_k^X(t) \\ &- i \sum_k \frac{\sqrt{\gamma_k}}{2} [C_{ab_k}^c - (C_{ab_k}^c)^* + C_{a^\dagger b_k}^c - (C_{a^\dagger b_k}^c)^*] dW_k^P(t) \end{aligned} \quad (24)$$

where the noted stochastic terms are visible in the second and third lines. These terms depend on joint cumulants of the cavity-QRC quantum state and are thus influenced by the QRC evolution; in particular, they can be nonzero even if Eq. (23) is satisfied. This principle of measurement-conditioned evolution has been analyzed in cQED experiments, for example in quantum trajectories of dispersively measured qubits nonreciprocally coupled to a quantum-limited phase-sensitive amplifier [68], and is accurately captured by the joint description of the measurement chain we employ here. In the present case, typical conditional cavity dynamics are illustrated for $\langle\hat{X}_a\rangle^c$ in Fig. 5(b), color-coded for the four states defined by $\Delta_a^{(\sigma)}$.

2. QRC

Our choice of coupling leads to the cavity \hat{X}_a quadrature driving the QRC via its \hat{X}_1 quadrature, whose resulting con-

ditional evolution is given by:

$$\begin{aligned} \frac{d}{dt}\langle\hat{X}_1\rangle^c &= \text{tr}\{\mathcal{L}_{\text{QRC}}\hat{\rho}^c\hat{X}_1\} - 2\Gamma_c\langle\hat{X}_a\rangle^c \\ &+ \sum_k \frac{\sqrt{\gamma_k}}{2} [C_{b_1 b_k}^c + (C_{b_1 b_k}^c)^* + 2C_{b_1^\dagger b_k}^c] dW_k^X(t) \\ &- i \sum_k \frac{\sqrt{\gamma_k}}{2} [C_{b_1 b_k}^c - (C_{b_1 b_k}^c)^*] dW_k^P(t) \end{aligned} \quad (25)$$

The QRC nodes thus evolve under the cavity signal, but also the QRC's intrinsic network dynamics governed by \mathcal{L}_{QRC} , and conditional evolution that depends on cumulants of the joint QRC state. Typical conditional dynamics of the $\langle\hat{X}_1\rangle^c$ quadrature of QRC node \hat{b}_1 for each generated cavity state is shown in Fig. 5(c). For this illustration, we have chosen two-node QRC parameters: $\gamma_1 = \gamma_2 = \kappa$, $\Lambda_1 = \Lambda_2 = 0.005\gamma_1$, $\Delta_1 = \Delta_2 = 0.0$, and $g_{12} = \gamma_1$.

Nonlinear processing of the cavity signal is encoded in this time evolution of the multimode QRC state. However, processed information about this state must be extracted from noisy measurement records $J_k^{X,P}(t)$ defined in Eqs. (4), (5) obtained from the measurement chain, which we discuss next.

B. Training and classification using finitely-sampled measurement distributions

Measurement records $\{J_k^{X,P}(t)\}$ contain experimentally-accessible information about the multimode QRC state for a given initialization of the measurement chain. We consider the case of *single-shot reservoir readout*, where the output $\mathbf{x}(t)$ of the QRC is constructed as the vector of all measured quadratures obtained via linear boxcar filtering of single-shot heterodyne measurement records [69],

$$\mathbf{x}(t) \equiv \begin{pmatrix} I_1^X(t) \\ I_1^P(t) \\ I_2^X(t) \\ I_2^P(t) \end{pmatrix}, \quad I_k^{X,P}(t) = \frac{1}{t} \int_0^t d\tau J_k^{X,P}(\tau). \quad (26)$$

Measured quadratures are themselves stochastic, so that the QRC output provides a distinct stochastic trajectory for repeated runs of the measurement chain; QRC output corresponding to $\Delta_a^{(\sigma)}$ is thus labelled $\mathbf{x}_{(q)}^{(\sigma)}$, indexed by both σ and a trajectory label q .

Having defined each element of the measurement chain and the measured QRC output, we can finally describe its operation for reservoir computing. We employ an output layer configured to yield a predicted class label σ^p using the measured QRC output, $\sigma^p = f_N\{\mathbf{W}_O \mathbf{x}_{(q)}^{(\sigma)} + \mathbf{b}\}$, where \mathbf{W}_O defines a matrix of time-independent trainable weights, \mathbf{b} a vector of trainable biases, and where $f_N\{\cdot\}$ is a fixed (that is, untrained) normalizing function. Perfect classification yields $\sigma^p = \sigma$ for all $\sigma = 1, \dots, C$, irrespective of q .

Formally, training is performed by minimizing a cost function, here the squared error over a training set whose output

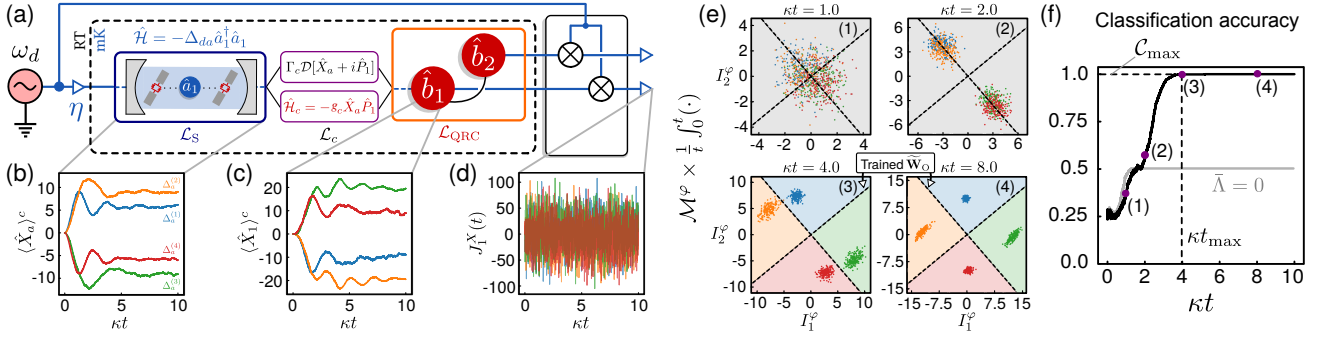


Figure 5. (a) Measurement chain using a 2-node QRC for the classification of states of a single cavity mode. The cavity mode is coupled to the QRC using a quantum-limited directional amplifier. Each node of the QRC is continuously monitored using separate output channels via heterodyne measurement. Conditional dynamics of (b) the cavity quadrature $\langle \hat{X}_a \rangle^c$, (c) the QRC node quadrature $\langle \hat{X}_1 \rangle^c$ it drives, and (d) measurement records $J_1^X(t)$, are shown for one initialization of the measurement chain per state, with colors blue, orange, green, and red corresponding to cavity mode detunings $\Delta_a^{(\sigma)} \in \{2.5, 1.5, -1.5, -2.5\}\kappa$ respectively. (e) Dynamics of QRC outputs projected onto the measured subspace $\{I_1^\varphi, I_2^\varphi\}$ as a function of time, with panels (1) through (4) corresponding to $\kappa t \in [1.0, 2.0, 4.0, 8.0]$ respectively. 100 trajectories are shown for each state σ . (f) Classification accuracy as a function of time, computed for a test set of size $Q_{\text{test}} = 200$ measured outputs per state. Black and gray curves correspond to a typical nonlinear QRC and a linear QRC respectively.

is known, and importantly is a convex minimization problem that is guaranteed to converge; full details of this training are included in Appendix J. Here we instead present a visual representation of training and classification, shedding light into how the QRC dynamics enable classification. Classification using linear weights \mathbf{W}_O is equivalent to constructing *linear* decision boundaries to separate reservoir outputs $\mathbf{x}_{(q)}^{(\sigma)}$ for different σ in the four-dimensional measured phase space spanned by $\mathbf{x} = (I_1^X, I_1^P, I_2^X, I_2^P)$, as shown in Appendix J. While this four-dimensional space is difficult to visualize, we find that for the present task two optimally chosen quadratures prove to be sufficient for successful classification. We thus rewrite the matrix of trainable weights as $\mathbf{W}_O \mathbf{x} = \widetilde{\mathbf{W}}_O \mathcal{M}^\varphi \mathbf{x}$, where \mathcal{M}^φ acts to project the measured QRC output onto a two-dimensional measured subspace $\{I_1^\varphi, I_2^\varphi\}$,

$$\begin{pmatrix} I_1^\varphi \\ I_2^\varphi \end{pmatrix} \equiv \mathcal{M}^\varphi \mathbf{x} = \begin{pmatrix} \cos \varphi_1 & \sin \varphi_1 & 0 & 0 \\ 0 & 0 & \cos \varphi_2 & \sin \varphi_2 \end{pmatrix} \begin{pmatrix} I_1^X \\ I_1^P \\ I_2^X \\ I_2^P \end{pmatrix}. \quad (27)$$

We note that all QRC node quadratures are still measured in each measurement run;. The projection phases are simply applied as part of the output layer, and are learned during training.

In Fig. 5(d), we project QRC outputs $\{\mathbf{x}_{(q)}^{(\sigma)}\}$ onto the measured subspace to visualize their evolution. Panels (1) through (4) show these outputs as a function of time, $\kappa t \in [1.0, 2.0, 4.0, 8.0]$ respectively. We see clearly that for short times (1), all four states are indistinguishable and the measured quadrature distributions overlap. This is not surprising, as the measurement chain is initialized to vacuum for all states. As the cavity is populated, distributions for states with distinct phases first separate (2) while the QRC response is still effectively linear. By (3), the cavity signal corresponding to higher amplitude states (orange, green) has increased suffi-

ciently to drive the QRC into a nonlinear regime. This is clear from the rotation of these distributions relative to those corresponding to the low amplitude states, as well as their visible distortion; the latter is a signature of amplified and squeezed QRC node fluctuations due to the Kerr nonlinearity, also seen in Fig. 3(c). The effect of the nonlinearity clearly separates all four distributions into (correspondingly shaded) distinct phase space regions.

It is then clear that linear decision boundaries (dashed black) can be drawn in the measured subspace to separate distributions for all C states. Training using a set of QRC output trajectories $\{\mathbf{x}_{(q)}^{(\sigma)}\}$ for $Q_{\text{Train}} = 100$ trajectories per state ($q = 1, \dots, Q_{\text{Train}}$) simply learns the parameters $\{\widetilde{\mathbf{W}}_O, \varphi\}$ uniquely specifying the optimal decision boundaries. Note that these decision boundaries, and therefore the training protocol to learn them, must take into account both the mean *and* variance of the finitely-sampled QRC output distributions. These quantities are affected by factors such as measurement time and filtering; therefore, training on measured output distributions is tailored to the finite experimental resources available for output processing. This is in sharp contrast to training on unconditional expectation values, which in principle assumes access to infinitely-many measurement records (for further discussion, see Appendix K). Finally, we emphasize that the linear output layer employed here serves two important, yet simple purposes: (i) filtering out measurement noise obscuring the underlying QRC dynamics, and (ii) drawing linear decision boundaries that optimally separate regions of the measured phase space into classes. The nontrivial processing step that enables classification - the evolution of measured outputs into linearly-separable regions of the measured subspace - is performed by the *nonlinear* QRC.

Once the QRC is trained, we can quantify its performance by predicting state labels σ^p for a distinct test set of reservoir outputs $\{\mathbf{x}_{(q)}^{(\sigma)}\}$ for $q = 1, \dots, Q_{\text{Test}} = 200$. Using the trained output layer, we determine σ^p for each $\mathbf{x}_{(q)}^{(\sigma)}$ and cal-

culate the classification accuracy as a function of time, defined as the fraction of $C \times Q_{\text{Test}}$ total trajectories that were correctly classified. The classification accuracy is plotted in Fig. 5(f), and is characterized by two metrics: the maximum classification accuracy C_{max} achieved within $\kappa t = 10.0$, and the time t_{max} taken to reach this maximum. We note that the specific QRC shown here achieves perfect classification ($C_{\text{max}} = 1$) before the measurement chain dynamics have settled into steady-state [see Figs. 5(b), (c)], indicating the ability of the QRC to process the information encoded in time-evolving quantum signals.

Also shown in gray is the classification accuracy for a linear QRC, $\bar{\Lambda} = 0$, which saturates to 0.5. The typical dynamics in the measured phase space are analogous to that shown in panel (2) of Fig. 5(e); the linear QRC is able to separate the pairs of states with different phases (blue-orange from red-green) but not distinguish states within this subset. Having used precisely the same form of output layer for both linear and nonlinear QRCs, we are able to unequivocally identify the QRC nonlinearity as being critical for the present classification task.

It is important to note, however, that a completely arbitrary nonlinear QRC will not succeed at this classification task; internal QRC parameters introduced in Eq. (2) must satisfy some task-dependent constraints, which we explore next.

C. Classification performance as a function of hyperparameters

Hyperparameters.— For large classical reservoirs, it is conventional - and practical - to describe reservoir parameters in terms of average hyperparameters instead of individual microscopic values. Here we adopt a similar convention despite the small size of the two-node QRC, parameterizing it as

$$\begin{aligned} \Lambda_k &= \bar{\Lambda} + [-\epsilon, \epsilon] \cdot \bar{\Lambda}, \quad \gamma_k = \bar{\gamma} = \kappa \quad \forall k, \\ \Delta_k &= \bar{\Delta} + [-\epsilon, \epsilon] \cdot \bar{\gamma}, \quad g_{12} = \bar{g} \end{aligned} \quad (28)$$

where $[x, y]$ denotes random samples from a uniform distribution with probability density between x and y ; we set $\epsilon = 0.1$ to observe the effect of random variations in QRC nonlinearities and frequencies that may occur in experiment. For convenience, the QRC damping rates are all set equal to the measured quantum system decay rate κ .

We analyze QRC classification performance as a function of the above hyperparameters in terms of the maximum classification accuracy C_{max} , and the time taken to reach this maximum, t_{max} . In doing so, we ensure that all other aspects of the measurement chain, training of the output layer, and testing remain unchanged from Sec. IV B, so that the particular role of QRC hyperparameters can be highlighted. Our key findings are recounted in a summary at the end of this section.

1. Dependence on nonlinearity $\bar{\Lambda}$

We begin by simulating QRCs with different values of the nonlinearity hyperparameter $\bar{\Lambda}$, while fixing $\bar{g}/\gamma = 1.0$,

$\bar{\Delta}/\bar{\gamma} = 0.0$, and calculating the maximum classification accuracy C_{max} . The results are shown in Fig. 6(a). Each red dot corresponds to a different random QRC, as defined by Eq. (28). The solid black curve is the mean over these QRCs, while the gray shaded region marks the range between worst and best performing QRCs. 5 random QRCs are simulated for each set of hyperparameters, amounting to over 200 distinct QRCs shown in Fig. 6(a).

For weak nonlinearities, C_{max} approaches 0.5, the limit of a linear QRC as seen earlier. With increasing nonlinearity however, the QRCs are able to perform the classification task, and C_{max} approaches unity. The time t_{max} to reach C_{max} can further distinguish operating hyperparameters: we plot κt_{max} in Fig. 6(a) (right-hand axis) in green, for parameters where $C_{\text{max}} \geq 0.99$ (yellow shaded region); green dots again indicate different random QRCs with the same hyperparameter values, while the solid green curve indicates the average, and the (small) green shaded region marks the range between the worst and best performing QRCs. We see that there exists a subset of optimal nonlinearity strengths for which the QRC correctly classifies all states in the shortest amount of time.

This calculation of C_{max} and t_{max} obtained from simulation of the complete measurement chain captures the full complexity of QRC evolution under time-dependent cavity signals with distinct amplitudes and phases. However, we find that a heuristic understanding of QRC performance characteristics can be developed by focusing on its *steady-state* response to a *time-independent* drive signal. For the pointer state classification task, we can define an effective drive strength $\eta_{\text{eff}}^{(\sigma)}$ as the long-time limit of the cavity signal corresponding to $\Delta_a^{(\sigma)}$ incident on the QRC,

$$\eta_{\text{eff}}^{(\sigma)} = -2\Gamma_c \lim_{t \rightarrow \infty} \{ \sqrt{2} \langle \hat{X}_a^{(\sigma)}(t) \rangle \} = -2\Gamma_c \eta \left[\frac{\Delta_a^{(\sigma)}}{(\Delta_a^{(\sigma)})^2 + \frac{\kappa^2}{4}} \right], \quad (29)$$

where the second equality is obtained from the steady-state solution of the unconditional version of Eq. (24). As visible in cavity signals in Fig. 5(b), for $|\Delta_a^{(\sigma)}| \geq \kappa$, $\sigma = 1, 4$ (blue, red) correspond to larger $|\Delta_a^{(\sigma)}|$ and yield weaker effective drives than $\sigma = 2, 3$ (orange, green).

We now analyze the steady-state QRC response in a regime enabling perfect classification, $\bar{\Lambda}/\bar{\gamma} = 0.005$ in Fig. 6(a) (and setting $\epsilon = 0$). To do so, we plot the steady-state values of $|\langle \hat{b}_{1,2} \rangle|$ (top panel) and $\arg \langle \hat{b}_{1,2} \rangle$ (bottom panel) under a constant drive strength η_{eff} in Fig. 6(b), with the particular effective drives $\eta_{\text{eff}}^{(\sigma)}$ indicated by correspondingly colored lines. Clearly, the QRC amplitude and phase response to higher amplitude cavity states $\sigma = 2, 3$ (orange, green) is non-linearly related to its response to lower amplitude states $\sigma = 1, 4$ (blue, red). This leads to the nonlinear amplitude response and phase rotation observed in the measured quadrature distributions for higher amplitude states in Fig. 5(e), which enables successful classification of the pointer states.

We find that the nonlinear QRC response observed in Fig. 6(b) is typical when the two-node QRC is successful at the pointer state classification task; it is therefore helpful to

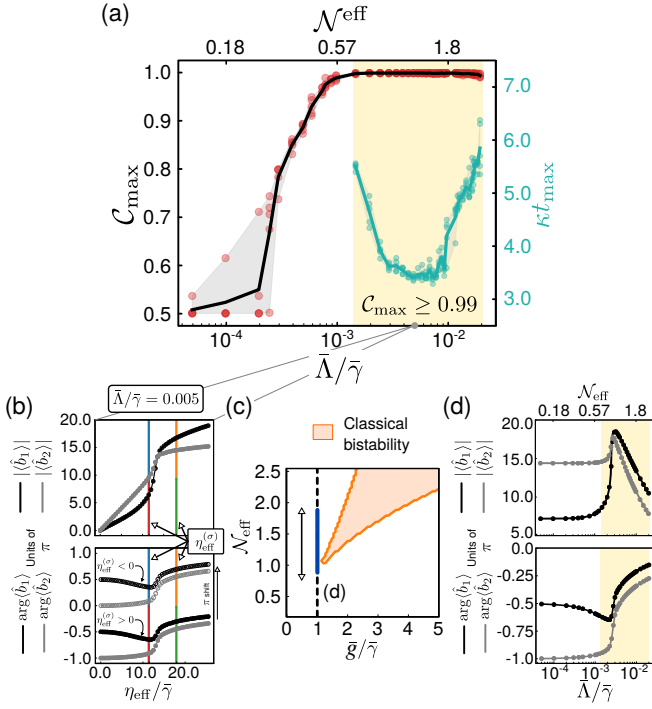


Figure 6. (a) C_{\max} (left-hand axis) vs. nonlinearity hyperparameter $\bar{\Lambda}$, for fixed coupling $\bar{g}/\bar{\gamma} = 1.0$ and detuning $\bar{\Delta}/\bar{\gamma} = 0.0$. For $C_{\max} \geq 0.99$ (yellow shaded region), κt_{\max} is plotted along the right-hand axis (green). Dots indicate different random QRCs, solid curves are averages over these QRCs, and shaded regions mark the range between worst and best performing QRCs (see text for more details). (b) Steady-state response of $|\langle \hat{b}_{1,2} \rangle|$, $\arg \langle \hat{b}_{1,2} \rangle$ as a function of constant drive strength η_{eff} . (c) Classical phase diagram of the two-node QRC in $(\bar{g}, \mathcal{N}_{\text{eff}})$ space, for fixed detuning $\bar{\Delta}/\bar{\gamma} = 0.0$. (d) Steady-state response of $|\langle \hat{b}_{1,2} \rangle|$ as a function of increasing nonlinearity $\bar{\Lambda}$ across the vertical dashed line ($\bar{g}/\bar{\gamma} = 1.0$) in (c).

locate such regimes in terms of general QRC hyperparameters. The *classical* phase diagram of the two-node QRC ends up providing the pathway to identifying such regimes. Furthermore, this phase diagram depends on the effective QRC drive and nonlinearity strengths only via the effective parameter (for details see Appendix G)

$$\mathcal{N}_{\text{eff}} = \frac{\eta_{\text{eff}}}{\bar{\gamma}} \sqrt{\frac{\bar{\Lambda}}{\bar{\gamma}}}. \quad (30)$$

To construct the QRC phase diagram informed by the pointer state classification task, we choose $\eta_{\text{eff}} \equiv \max_{\sigma} \{|\eta_{\text{eff}}^{(\sigma)}|\}$, corresponding to higher amplitude cavity states that drive the QRC into nonlinear processing regimes.

We plot the resulting phase diagram in Fig. 6(c) as a function of \mathcal{N}_{eff} and coupling \bar{g} , for fixed $\bar{\Delta}/\bar{\gamma} = 0.0$. The orange region indicates the classical bistability, analogous to that found for the single QRC node. The vertical dashed line indicates $\bar{g}/\bar{\gamma} = 1.0$, namely the cross-section along which Fig. 6(a) is plotted. Also shown in Fig. 6(d) is the steady-state response of the QRC node amplitudes and phases along this cross-section; the solid blue line in (c) and yellow shaded

region in (d) mark \mathcal{N}_{eff} values for which $C_{\max} \geq 0.99$. For weak nonlinearities where classification is unsuccessful, the mode amplitudes and phases are effectively constant, indicating the linear regime of operation where QRC states generated by different η_{eff} cannot be separated. For increasing $\bar{\Lambda}$, approaching the classical bistability results in a nonlinear QRC response enabling successful classification. When the nonlinearity is further increased however, $|\langle \hat{b}_{1,2} \rangle|$ begins to decrease; this is a result of the increasing Kerr shift moving the node resonances further from the steady-state input signals. The reduced QRC amplitude is responsible for the increase in t_{\max} seen in Fig. 6(a).

2. Dependence on coupling \bar{g} and detuning $\bar{\Delta}$

We can similarly analyze the performance of QRCs as a function of coupling and detuning hyperparameters, for fixed nonlinearity (here, $\bar{\Lambda}/\bar{\gamma} = 0.002$, which fixes $\mathcal{N}_{\text{eff}} = 0.8$). In cQED realizations of the proposed QRCs, these hyperparameters can be rendered *in-situ* tunable (via flux-modulated SQUIDs and parametric couplers respectively), thus providing experimental control over QRC performance.

We begin by calculating C_{\max} as a function of the coupling hyperparameter \bar{g} in Fig. 7(a), fixing the detuning $\bar{\Delta}/\bar{\gamma} = 0.0$. This cross-section of $(\bar{\Delta}, \bar{g})$ space is indicated in the center panel classical phase diagram with a vertical dashed line, with the corresponding QRC steady-state amplitude and phase response plotted in the left panel of Fig. 7(c). For weak couplings, the \hat{b}_2 is correspondingly weakly excited, since the input couples only to \hat{b}_1 . Only the \hat{b}_1 is effectively available for processing; this is insufficient for the task at hand and the QRC does not attain $C_{\max} = 1$ within the measurement time.

With increasing coupling, the QRC achieves perfect classification; furthermore, an optimal range of couplings exists where t_{\max} (green, right-hand axis) is shortest. From the left panel of Fig. 7(c), we see that here the QRC response for both nodes is strongest. Curiously, the performance dips again for strong couplings: here the dynamics due to the linear coupling dominate over the nonlinearity, as the QRC polariton modes are further split and respond weakly to the signal emanating from the cavity, resulting in the decreased response visible in the steady state plots.

To demonstrate the advantage of *in-situ* control, we now consider QRCs with sub-optimal coupling $\bar{g}/\bar{\gamma} = 2.0$ (where $C_{\max} \simeq 0.8$, see Fig. 7(a)). Figure 7(b) shows the resulting C_{\max} as a function of the detuning hyperparameter $\bar{\Delta}$, a range indicated via the horizontal dashed line in the phase diagram of Fig. 7(c). The corresponding steady-state QRC amplitude and phase responses are shown in the right panel. We clearly see that C_{\max} approaches unity once more for specific positive and negative detunings where the QRC response is enhanced.

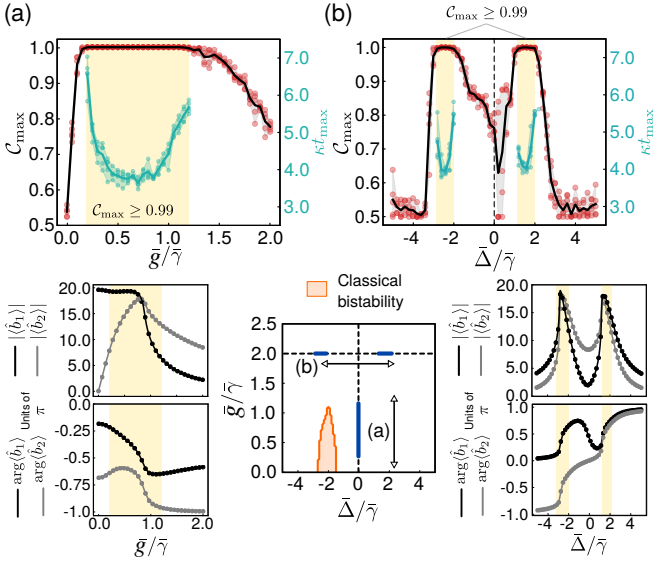


Figure 7. C_{\max} (left-hand axis) for fixed nonlinearity $\bar{\Lambda}/\bar{\gamma} = 0.002$ (and hence fixed $\mathcal{N}_{\text{eff}} = 0.8$) vs. (a) coupling hyperparameter \bar{g} for fixed detuning $\bar{\Delta}/\bar{\gamma} = 0.0$, and vs. (b) average detuning hyperparameter $\bar{\Delta}$ for fixed coupling $\bar{g}/\bar{\gamma} = 2.0$. For both plots, where $C_{\max} \geq 0.99$ (yellow shaded region), t_{\max} is plotted along the right-hand axis (green). Dots indicate different random QRCs, solid curves are averages over these QRCs, and shaded regions mark the range between worst and best performing QRCs (see text for more details). (c) Classical phase diagram of the two-node QRC in $(\bar{\Delta}, \bar{g})$ space for fixed $\mathcal{N}_{\text{eff}} = 0.8$. Cross-sections of the steady-state response of QRC nodes are shown in (d) for fixed detuning $\bar{\Delta}/\bar{\gamma} = 0.0$ as a function of coupling \bar{g} , and in (e) for fixed coupling $\bar{g}/\bar{\gamma} = 2.0$ as a function of detuning $\bar{\Delta}$.

3. Dependence on damping $\bar{\gamma}$

Thus far, the QRC node damping rate has appeared as the natural scaling parameter for hyperparameters in our analysis. Hence if $\bar{\gamma}$ is varied but the remaining QRC hyperparameters are suitably adjusted so that $\{\mathcal{N}_{\text{eff}}, \bar{g}/\bar{\gamma}, \bar{\Delta}/\bar{\gamma}\}$ are unchanged, the QRC will operate in a fixed regime of the classical phase diagram; provided this regime enables nonlinear processing as previously analyzed, successful classification will be possible.

However, two aspects of QRC dynamics controlled by $\bar{\gamma}$ are not captured by its steady-state phase diagram: the response time of QRC nodes to *time-dependent* input signals, as well as the magnitude of QRC response for fixed input signal strengths. To isolate the influence of $\bar{\gamma}$ on QRC performance via these specific factors, we must ensure that the QRC operates in a fixed nonlinear regime of the classical phase diagram even as $\bar{\gamma}$ is varied.

This analysis is carried out in Appendix H, with our main results summarized here. We find that for $\bar{\gamma} \ll \kappa$, classification is still possible in nonlinear QRC regimes; however, as the QRC responds much slower than the evolution timescale of cavity signals (set by κ), t_{\max} for classification increases. As $\bar{\gamma}$ increases the QRC response time is reduced, leading to a decrease in t_{\max} until a minimum is attained around $\bar{\gamma} \simeq 2\kappa$. For larger $\bar{\gamma}$, while the QRC response time is still fast rela-

tive to the quantum system evolution timescale, the QRC node amplitude is suppressed such that the amplitude of extracted output signals from the measurement chain is reduced relative to the measurement noise, thereby increasing t_{\max} once more.

4. Summary

The results of this section indicate that a range of hyperparameter values $\{\mathcal{N}_{\text{eff}}, \bar{g}/\bar{\gamma}, \bar{\Delta}/\bar{\gamma}\}$ exist for which the QRC is able to perform nonlinear processing of input signals, and consequently succeed at the task of pointer state classification considered here. Importantly, regions of classical bistability - key signatures of classical nonlinearity - are strong (although not exclusive) indicators of such parameter regimes. The mechanism of classification is the nonlinear response of QRC node amplitudes with respect to the amplitude and phase of incident signals from the cavity mode.

D. Dependence on input signal strength: classical vs. quantum regimes of operation

Our analysis of QRC performance as a function of hyperparameters highlights the importance of operating in dynamical regimes that enable nonlinear processing of input signals. Importantly, the effective nonlinearity parameter \mathcal{N}_{eff} that is central to such regimes depends not only on the QRC nonlinearity $\bar{\Lambda}$, but also on the effective drive η_{eff} . The latter is ultimately related to the strength of the drive η incident on the measurement chain, which in turn can be constrained by various physical considerations. For composite measurement chains considered here, the incident field initially drives the measured quantum system upstream of the QRC. Hence, its amplitude is dictated by the desired response of this quantum system, namely the amplitude required to generate quantum states of interest, as well as to avoid possible non-idealities that may appear under strong driving. Secondly, to decrease the energy of physical computation, one must naturally consider a reduction in the input power to the measurement chain containing the QRC.

With this view, we consider separate instances of pointer state classification tasks for the same values of detunings $\Delta_a^{(\sigma)}$ defining the pointer states as considered earlier, but for different values of drive input to the measurement chain, $\eta/\bar{\gamma} \in \{26.0, 15.0, 9.0\}$. These input drive values translate to effective QRC drives $\eta_{\text{eff}}/\bar{\gamma} \in \{31.2, 18.0, 10.8\}$. We then determine classification metrics C_{\max} and κt_{\max} as a function of QRC nonlinearity strength $\bar{\Lambda}$ as before, holding the remaining QRC hyperparameters fixed, $\bar{g}/\bar{\gamma} = 1.0$, $\bar{\Delta}/\bar{\gamma} = 0.0$.

The achieved C_{\max} is plotted in the top panel of Fig. 8 for the aforementioned $\eta_{\text{eff}}/\bar{\gamma}$ values (blue, purple, and peach respectively). Immediately, we note that for weaker input drives, a larger QRC nonlinearity $\bar{\Lambda}$ is required for successful pointer state classification. When interpreted in the context of the effective nonlinearity \mathcal{N}_{eff} , this result makes sense: the latter depends on the product of the nonlinearity $\bar{\Lambda}$ of the QRC,




Figure 8. C_{\max} (top panel) and t_{\max} (bottom panel) as a function of average nonlinearity hyperparameter $\bar{\Lambda}$, for instances of the pointer state classification task corresponding to three different effective QRC drives $\eta_{\text{eff}}/\bar{\gamma}$. Dots indicate different random QRCs, solid curves are averages over these QRCs, and shaded regions mark the range between worst and best performing QRCs. Crosses indicate QRCs with fixed $\mathcal{N}_{\text{eff}} \simeq 1.4$ for the three considered task instances; for weaker $\eta_{\text{eff}}/\bar{\gamma}$, corresponding QRCs have stronger nonlinearity and operate in more quantum regimes (see text for more details).

and the effective drive η_{eff} it sees. To operate QRCs in a nonlinear processing regime given by a fixed value of \mathcal{N}_{eff} (for example, $\mathcal{N}_{\text{eff}} \simeq 1.4$ marked by colored crosses in Fig. 8), a weaker η_{eff} will require stronger $\bar{\Lambda}$.

This intuitive result encourages a complementary perspective on reservoir processing across varying input drive strengths, enabled by our full quantum treatment of the measurement chain using truncated cumulants. As shown in Sec. III, increasing $\bar{\Lambda}$ for fixed \mathcal{N}_{eff} controllably transitions Kerr QRCs from classical to quantum regimes of operation. Therefore, considering instances of the pointer state classification task across drive strengths allows us to analyze and compare QRC performance on the task throughout this transition.

In this context, extrapolating the results of Fig. 8 to weaker nonlinearities leads to an immediate observation: even a QRC operating in the classical nonlinear regime ($\bar{\Lambda} \rightarrow 0$ for fixed \mathcal{N}_{eff}) can perform the pointer state classification task we consider here, albeit requiring increasing input strengths. In this limit, the quantum state of the measurement chain is described as a product of coherent states, defined entirely by its first-

order cumulants. Fortunately, the information that distinguishes qubit pointer states is also encoded in their first-order cumulants (amplitudes and phases) and not their second-order cumulants (quantum fluctuations). As a result, the nonlinear mapping of this information to measured outputs by the nonlinear QRC is described primarily by dynamics of first-order cumulants, and is therefore possible in the classical limit of QRC operation.

Fig. 8 indicates that for weaker incident drive strengths, classification is enabled by QRCs with stronger nonlinearities, operating in more quantum regimes; however, not all classification metrics are unchanged. In particular, the time t_{\max} to reach maximum classification accuracy (plotted in the lower panel of Fig. 8) *increases* for weaker input drives. This can be understood as follows. Decreasing the input drive strength reduces the amplitude (and hence first-order cumulants) of the pointer states to be classified. This reduced amplitude can be compensated by increasing the QRC nonlinearity to perform the desired nonlinear mapping to measured outputs; however, the amplitude of QRC outputs is still correspondingly reduced. In terms of the measured subspace $\{I_1^\varphi, I_2^\varphi\}$ shown in Fig. 5(e), the QRC output distributions to be separated have centroids that are closer to each other. The spread of these distributions, on the other hand, is determined by measurement noise (independent of drive strength) and QRC-mediated amplified or squeezed quantum fluctuations (mostly unchanged for fixed \mathcal{N}_{eff}); it primarily decreases with integration time t due to filtering. Hence a decrease in input strengths reduces the relative separation of the distributions for a given measurement time, making them harder to separate, and demanding longer integration times t_{\max} to enable classification.

The above results have an important consequence for the analysis of quantum reservoir processing for general computational tasks. Information processing with reduced input signals necessitates strongly-nonlinear QRCs, and thus provides a physically-motivated probe of QRC operation in quantum regimes. However, by using a framework that accounts for the complete measurement chain, we explicitly observe that the decrease in input power can lead to an increase in required measurement resources for computation; in this particular case, the measurement time. Such costs must be accounted for in any analysis of the possible advantages of reservoir computing in the quantum regime.

Finally, we note that processing with low input signals is commonplace in cQED experiments. Constraints on measurement time due to qubit decay and the use of noisy HEMT amplifiers [70] in cQED experiments means that the reduced output signal is overcome by including quantum-limited amplifiers [71–73] as part of the measurement chain. Such amplifiers can be included as part of our QRC framework for operation with low signal powers, as additional quantum measurement resources. The key considerations in doing so are discussed in Appendix L (albeit applied to our second computational task, which is considered next).

V. AMPLIFIER MEASUREMENT: CLASSIFYING GAUSSIAN STATES WITH EQUAL MEANS AND DISTINCT VARIANCES

We will now apply our quantum reservoir computing framework to Task II: the classification of Gaussian states that differ only in their second-order moments, or variances. To generate such states, we consider a minimal system based on a two-mode amplifier; as a shorthand, we refer to this task as that of amplifier state classification from here on.

A. Hardware measurement chain

The quantum measurement chain is depicted schematically in Fig. 9(a). The quantum system is coupled to a single-node QRC, which we will see proves sufficient for the classification task at hand. We again begin with a description of each element of this measurement chain, depicted in Fig. 9(a).

1. Two-mode amplifier and nonreciprocal coupling to QRC

Our measured quantum system is a general two-mode amplifier, described by the Liouvillian superoperator \mathcal{L}_S defined as

$$\mathcal{L}_S \hat{\rho} = -i[\hat{\mathcal{H}}_0 + \hat{\mathcal{H}}_S^{(\sigma)} + \eta^{(\sigma)}(-i\hat{a}_1 + i\hat{a}_1^\dagger), \hat{\rho}] + \sum_n \kappa_n \mathcal{D}[\hat{a}_n] \hat{\rho}, \quad (31)$$

where the system Hamiltonians take the form:

$$\begin{aligned} \hat{\mathcal{H}}_0 &= - \sum_{n=1,2} \Delta_{an} \hat{a}_n^\dagger \hat{a}_n, \\ \hat{\mathcal{H}}_S^{(\sigma)} &= \frac{1}{2} G_1^{(\sigma)} e^{-i\pi/2} \hat{a}_1^2 + G_{12}^{(\sigma)} \hat{a}_1 \hat{a}_2 + h.c., \end{aligned} \quad (32)$$

where $\Delta_{an} = \omega_d - \omega_{an}$. The two-mode system is defined by a degenerate parametric amplifier interaction $\propto G_1$ acting on mode \hat{a}_1 , and a non-degenerate interaction $\propto G_{12}$. Note that these interaction strengths can also be set by tunable parametric drives in cQED implementations. Defining for convenience $\hat{X}_{a_1} = \frac{1}{\sqrt{2}}(\hat{a}_1 + \hat{a}_1^\dagger)$, $\hat{P}_{a_1} = -\frac{i}{\sqrt{2}}(\hat{a}_1 - \hat{a}_1^\dagger)$ as the canonically conjugate quadratures of the \hat{a}_1 mode, note that our choice of drive Hamiltonian $\propto \hat{P}_{a_1}$ drives the amplifier \hat{X}_{a_1} quadrature only. Distinct states of this system are thus defined by the drive strength and values of the two interaction parameters, $(\eta^{(\sigma)}, G_1^{(\sigma)}, G_{12}^{(\sigma)})$.

As before, understanding the dynamics of this two-mode amplifier requires specifying how it is embedded within the measurement chain including the QRC. We choose to couple the system \hat{a}_1 mode to the \hat{b}_1 node of the QRC; distinct from Sec. IV, however, we now employ a non-reciprocal hopping interaction (modeling a standard circulator), defined by the coupling superoperator \mathcal{L}_c ,

$$\mathcal{L}_c \hat{\rho} = -i \left[\frac{i}{2} g_c \hat{a}_1^\dagger \hat{b}_1 + h.c., \hat{\rho} \right] + \Gamma_c \mathcal{D}[\hat{a}_1 + \hat{b}_1] \hat{\rho}. \quad (33)$$

To understand the role of this coupling, we can once again analyze the conditional dynamics of the \hat{X}_{a_1} quadrature that couples to the QRC:

$$\begin{aligned} \frac{d}{dt} \langle \hat{X}_{a_1} \rangle^c &= \text{tr}\{\mathcal{L}_S \hat{\rho}^c \hat{X}_{a_1}\} + \frac{g_c - \Gamma_c}{2} \langle \hat{X}_1 \rangle^c - \frac{\Gamma_c}{2} \langle \hat{X}_{a_1} \rangle^c \\ &+ \frac{\sqrt{\gamma_1}}{2} \left[C_{a_1 b_1}^c + (C_{a_1 b_1}^c)^* + C_{a_1^\dagger b_1}^c + (C_{a_1^\dagger b_1}^c)^* \right] dW_1^X(t) \\ &- i \frac{\sqrt{\gamma_1}}{2} \left[C_{a_1 b_1}^c - (C_{a_1 b_1}^c)^* + C_{a_1^\dagger b_1}^c - (C_{a_1^\dagger b_1}^c)^* \right] dW_1^P(t) \end{aligned} \quad (34)$$

Setting $g_c = \Gamma_c$ once again ensures that the amplifier does not see the QRC field, while the QRC is driven by the field from mode \hat{a}_1 . However, in addition to this coupling term, the form of \mathcal{L}_c introduces local damping $\propto \Gamma_c$ of the \hat{a}_1 mode, which modifies the system's dynamics. For convenience, we choose damping parameters such that the total damping rate of system modes \hat{a}_1 and \hat{a}_2 is equal after taking this additional damping term into account; more precisely, we set $\kappa_1 + \Gamma_c = \kappa_2 \equiv \kappa$. For simplicity, we further set $\kappa_1 = \Gamma_c = \kappa/2$, and $\omega_d = \omega_{a_1} = \omega_{a_2}$ so that the detunings $\Delta_{a_1} = \Delta_{a_2} = 0$.

We are now in a position to define our classification task: we wish to use a QRC to classify $C = 2$ different quantum states of the two-mode amplifier, defined by specific values of the drive and amplifier interaction strengths $(\eta^{(\sigma)}, G_1^{(\sigma)}, G_{12}^{(\sigma)}) \in \{(5.0, 0.3, 0.0), (8.0, 0.0, 0.3)\} \kappa$. Leaving details for Appendix I, these parameter choices generate amplifier states with $\langle \hat{X}_{a_1}^{(1)} \rangle = \langle \hat{X}_{a_1}^{(2)} \rangle \neq 0$, and $\langle \hat{P}_{a_1}^{(1)} \rangle = \langle \hat{P}_{a_1}^{(2)} \rangle = 0$, namely with equal first-order moments for the mode \hat{a}_1 that the QRC is directly coupled to. This is visible in the conditional dynamics of $\langle \hat{X}_{a_1} \rangle^c$ plotted in Fig. 9(b), with blue and orange colors corresponding to the two states. However the two states are distinct, as is clear from their different second-order cumulants, whose unconditional steady-state values are shown in the bar plots in Fig. 9(b).

2. QRC

Under directional coupling $g_c = \Gamma_c$, the amplifier \hat{a}_1 mode drives the QRC via the \hat{b}_1 mode. For specificity, we can look at the conditional evolution of the \hat{X}_1 quadrature of the QRC:

$$\begin{aligned} \frac{d}{dt} \langle \hat{X}_1 \rangle^c &= \text{tr}\{\mathcal{L}_{\text{QRC}} \hat{\rho}^c \hat{X}_1\} - \Gamma_c \langle \hat{X}_{a_1} \rangle^c - \frac{1}{2} \Gamma_c \langle \hat{X}_1 \rangle^c \\ &+ \frac{\sqrt{\gamma_1}}{2} \left[C_{b_1 b_1}^c + (C_{b_1 b_1}^c)^* + 2C_{b_1^\dagger b_1}^c \right] dW_1^X(t) \\ &- i \frac{\sqrt{\gamma_1}}{2} \left[C_{b_1 b_1}^c - (C_{b_1 b_1}^c)^* \right] dW_1^P(t) \end{aligned} \quad (35)$$

The QRC nodes thus evolve under the amplifier signal and also experiences local damping due to the coupling superoperator, in addition to its internal dynamics governed by \mathcal{L}_{QRC} , and conditional evolution that depends on cumulants of the QRC state. Typical conditional dynamics of $\langle \hat{X}_1 \rangle^c$ are shown in Fig. 5(c) for the two amplifier states. The specific single-node QRC parameters we have chosen here are $\gamma_1 = \kappa$, $\Lambda_1 = 0.0027\gamma_1$, and $\Delta_1 = -\gamma_1$.

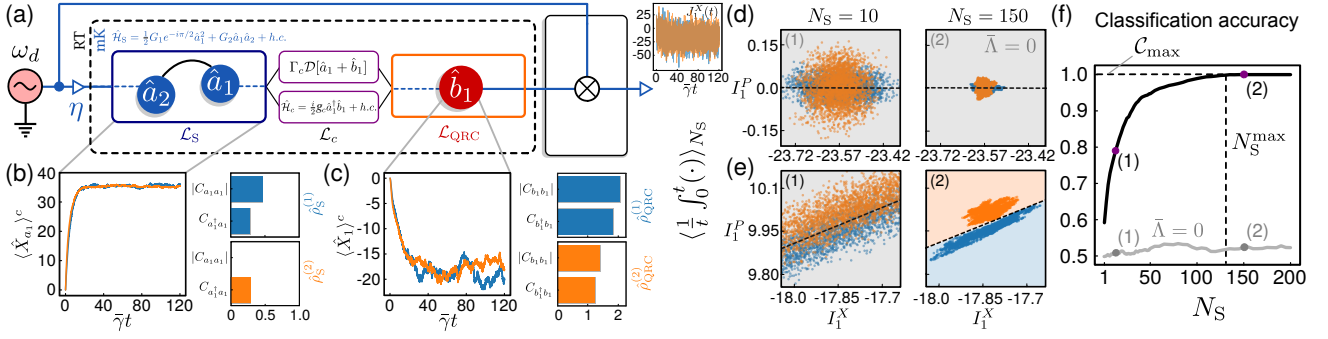


Figure 9. (a) Measurement chain using a single-node QRC for the classification of Gaussian states generated by a two-mode amplifier. A single mode \hat{a}_1 of the system is coupled to the QRC using a directional hopping interaction (a circulator). The single QRC node is continuously monitored via heterodyne measurement. (b) Conditional dynamics of the \hat{a}_1 mode's $\langle \hat{X}_a \rangle^c$ quadrature, and steady-state unconditional cumulants. (c) Same as (b), for the \hat{b}_1 QRC node's $\langle \hat{X}_1 \rangle^c$ quadrature and unconditional cumulants. Conditional dynamics are shown for one initialization of the measurement chain per state. Colors blue and orange correspond to parameter sets $(\eta^{(\sigma)}, G_1^{(\sigma)}, G_{12}^{(\sigma)}) \in \{(5.0, 0.3, 0.0), (8.0, 0.0, 0.3)\} \kappa$ respectively. (d) Measured phase space quadratures $\{I_1^X(t_f), I_1^P(t_f)\}$ for (1) $N_S = 10$, and (2) $N_S = 150$, for a linear QRC, $\bar{\Lambda} = 0.0$. (e) Same as (d) but for a nonlinear QRC (see text for parameters). (f) Classification accuracy as a function of N_S , computed for a test set of size $Q_{\text{Test}} = 3000$; gray and black curves correspond to the linear and nonlinear QRCs of (d) and (e) respectively.

The QRC is driven to distinct quantum states under the two amplifier state inputs, as is clear from the unconditional second-order QRC cumulants shown in the bar plots in Fig. 5(c). What is not immediately clear is that the resulting QRC states can *also* differ in their first-order moments, even though the two amplifier states do not; this has implications for state classification, as we analyze next.

B. Training and classification using finitely-sampled measurement distributions

Measurement records $\{J_1^{X,P}(t)\}$ obtained from the measurement chain contain information about QRC dynamics that will enable state classification. In using these measurement records to construct the measured QRC output, we will consider a slightly different scheme that is also experimentally relevant, to highlight the flexibility afforded by our description.

In contrast to the pointer state classification task, the measured QRC output $\mathbf{x}(t)$ here is constructed as the vector of measured quadratures obtained using both temporal filtering *and* ensemble-averaging of heterodyne measurement records, calculated over N_S distinct runs of the measurement chain, or ‘shots’,

$$\mathbf{x}(t) \equiv \begin{pmatrix} I_1^X(t) \\ I_1^P(t) \end{pmatrix}, \quad I_1^{X,P}(t) = \left\langle \frac{1}{t} \int_{t_0}^t d\tau J_1^{X,P}(\tau) \right\rangle_{N_S}. \quad (36)$$

This more general form of readout (which reduces to the single-shot case when $N_S = 1$) can be used to reduce noise power in readout using multiple measurement records obtained for the same inputs to the measurement chain. Measured quadratures are themselves stochastic, so that the QRC output provides a distinct stochastic trajectory for repeated runs of the measurement chain. QRC output corresponding to

different parameters $\{\eta^{(\sigma)}, G_1^{(\sigma)}, G_{12}^{(\sigma)}\}$ is thus labelled $\mathbf{x}_{(q)}^{(\sigma)}$, where q is the trajectory label. Note that each trajectory q for a given σ is constructed using N_S measurement records, each obtained from a distinct run of the measurement chain for the same input parameters $\{\eta^{(\sigma)}, G_1^{(\sigma)}, G_{12}^{(\sigma)}\}$.

Secondly, instead of integrating the entire obtained measurement records, we introduce a waiting time t_0 . By choosing t_0 to be much longer than the amplifier response time for the two states (specifically, $t_0 = 35.0/\kappa \gg \frac{2}{\kappa - 2G_1^{(1)}} \simeq \frac{2\kappa}{\kappa^2 - 4(G_{12}^{(2)})^2}$), we ensure that initial transients have settled, so that any information about the states contained in their slightly-distinct transient dynamics is lost and cannot influence classification.

The classification task is then set up as in Sec. IV. We employ an output layer to predict a class label σ^p using the measured QRC output, $\sigma^p = f_N\{\mathbf{W}_O \mathbf{x}_{(q)}^{(\sigma)}(t_f) + \mathbf{b}\}$. Here, \mathbf{W}_O defines a matrix of time-independent trainable weights, \mathbf{b} a vector of trainable biases, while $f_N\{\cdot\}$ is the same fixed (that is, untrained) normalizing function from earlier. Perfect classification yields $\sigma^p = \sigma$ for all $\sigma = 1, \dots, C$. Note that here we consider the measured phase space as being spanned by measured quadratures evaluated at the *end* of the measurement, namely $\mathbf{x}(t_f) = \{I_1^X(t_f), I_1^P(t_f)\}$ with $t_f = 120.0/\kappa$. This corresponds to classification using a fixed measurement time; we will instead analyze performance as a function of the other measurement resource, the number of shots N_S .

Beginning with the case of a complete linear QRC ($\bar{\Lambda} = 0$), the QRC outputs $\{\mathbf{x}_{(q)}^{(\sigma)}\}$ are shown in measured phase space in Fig. 9(d), for $q = 1, \dots, 1500$ per state σ (blue and orange). Each panel shows outputs for the indicated value of N_S . We see (1) that for $N_S = 10$, the distributions corresponding to the two states have distinct profiles but are overlapping, and hence not separable by a linear decision boundary (black dashed line). The spread of the distributions can be decreased by averaging over an increasing number of shots;

this is seen in (2), where $N_S = 150$ (the phase space area shown is the same as in (1), allowing direct comparison of the distribution area). However, the centers of the measured distributions are unshifted, so that they still overlap and their linear separation is not possible.

The observed distributions are very different for an appropriately-chosen nonlinear QRC. The measured phase space is shown in Fig. 9(e), again for increasing values of N_S . For (1) $N_S = 10$, we see that the distributions are significantly distorted in comparison to the linear QRC, showing amplified and squeezed axes due to the Kerr nonlinearity. However the distributions still overlap significantly and are not linearly separable. With increasing (2) $N_S = 150$, the width of the distributions decreases, and a relative displacement of their centers (means) ensures that they can be separated by the dashed black linear decision boundary, allowing for perfect classification.

Note that both the linear and nonlinear QRC's are driven to distinct quantum states under input from the two states to be classified; this is clear from the distinct distributions observed in the measured phase space. However, that this difference manifests in a change to first order moments that are linearly related to the measured quadratures is a signature of nonlinear quantum systems, and not specific to the Kerr model we have chosen here: first-order moments can couple directly to higher-order moments, and can thus be modified by differences in these higher-order moments. Precisely this effect was observed in the steady-state response of a single driven Kerr QRC node in Figs. 3(b), (c).

As before, an arbitrary nonlinear QRC will not generally carry out this classification task optimally. To understand the required operating regimes, it proves useful to analyze the classification performance more quantitatively. Using a training set $\{\mathbf{x}_{(q)}^{(\sigma)}\}$ for $q = 1, \dots, Q_{\text{Train}} = 1500$ QRC outputs per state, we learn the optimal weights \mathbf{W}_O and biases \mathbf{b} defining the linear decision boundary separating measured output distributions. Then, the trained QRC is used to predict state labels σ^p for a distinct test set $\{\mathbf{x}_{(q)}^{(\sigma)}\}$ for $q = 1, \dots, Q_{\text{Test}} = 3000$, and the classification accuracy as a function of N_S is calculated. The result is plotted in Fig. 9(f), in gray and black for the linear and the specific nonlinear QRC respectively, and is characterized by two metrics: the maximum classification accuracy \mathcal{C}_{max} as before, and the number of shots used in constructing measured quadratures to reach this maximum, N_S^{max} . In the following section, we will analyze these performance metrics as a function of QRC hyperparameters to determine constraints that enable the amplifier state classification task.

C. Classification performance as a function of hyperparameters

Hyperparameters.— For a single-node QRC, the parameter space to vary is reduced relative to the task considered in Sec. IV, which employed a two-node QRC. Regardless, for consistency with our prior analysis we employ a similar notation, setting $\Lambda_1 \equiv \bar{\Lambda}$, $\Delta_1 \equiv \bar{\Delta}$, and setting the damping

rate equal to the (common) decay rate κ of amplifier modes, $\gamma_1 \equiv \bar{\gamma} = \kappa$. Once again, all remaining aspects of the measurement chain and training of the output layer remain unchanged as hyperparameters are varied and QRC performance observed. Our key findings are summarized at the end of this section.

1. Dependence on nonlinearity $\bar{\Lambda}$

We begin by fixing the QRC detuning $\bar{\Delta}/\bar{\gamma} = -1.0$, and calculating \mathcal{C}_{max} for varying QRC nonlinearity $\bar{\Lambda}$; the results are shown in Fig. 10(a). Clearly, an optimal range of nonlinearity strengths exists where $\mathcal{C}_{\text{max}} \rightarrow 1$. For the orange shaded region where $\mathcal{C}_{\text{max}} \geq 0.99$, we also plot N_S^{max} along the right-hand axis (green curve). A nonlinearity strength exists for which N_S^{max} is minimized, thereby reducing the measurement resources required for classification.

The binary nature of the current task allows us to extract further intuition about the observed QRC performance. As discussed earlier, the ability to linearly separate integrated output records processed by the QRC corresponding to the two amplifier states depends on the measured noisy distributions acquiring a sufficient displacement. The unconditional QRC state variable that is a measure of this displacement is simply the difference in QRC node amplitude for the two states, $\mathcal{B}^{12} \equiv |\langle \hat{b}_1^{(1)} \rangle - \langle \hat{b}_1^{(2)} \rangle|$. We plot this quantity calculated using TEOMs in the top panel of Fig. 10(b) for the same parameter regime as (a), together with the yellow shaded region where $\mathcal{C}_{\text{max}} \geq 0.99$. The lower panel shows unconditional cumulants of the QRC state corresponding to the two amplifier states (blue and orange respectively). Outside the yellow shaded region, the QRC state is driven to distinct cumulants for the two amplifier states, but the magnitude of cumulants, and the resulting displacement of QRC node amplitudes \mathcal{B}^{12} , is small. However, within the yellow shaded region, the distinct cumulants are much larger, leading to larger value of \mathcal{B}^{12} . In contrast, for a linear QRC (dashed gray line), \mathcal{B}^{12} is always zero.

We emphasize here that a nonzero value of this displacement is by itself insufficient for successful discrimination between the amplifier states. The displacement is contained within noisy measurement records, and classification amounts to being able to distinguish their filtered, ensemble-averaged distributions in measured phase space. These factors are all accounted for by our training on measured quadrature distributions, which yield all our results for \mathcal{C}_{max} and N_S^{max} .

Operating regimes where \mathcal{B}^{12} becomes nonzero can once again be analyzed in the context of the classical phase diagram of the single-node QRC. To this end, we introduce an effective drive η_{eff} experienced by the QRC due to its coupling to the amplifier mode \hat{a}_1 . Unlike the QND interaction employed in Sec. IV which coupled only the cavity \hat{X}_a quadrature to the QRC, the directional hopping interaction here ensures that mode \hat{b}_1 sees both quadratures of mode \hat{a}_1 . The effective drive $\eta_{\text{eff}}^{(\sigma)}$ seen by the QRC for each amplifier state then takes the

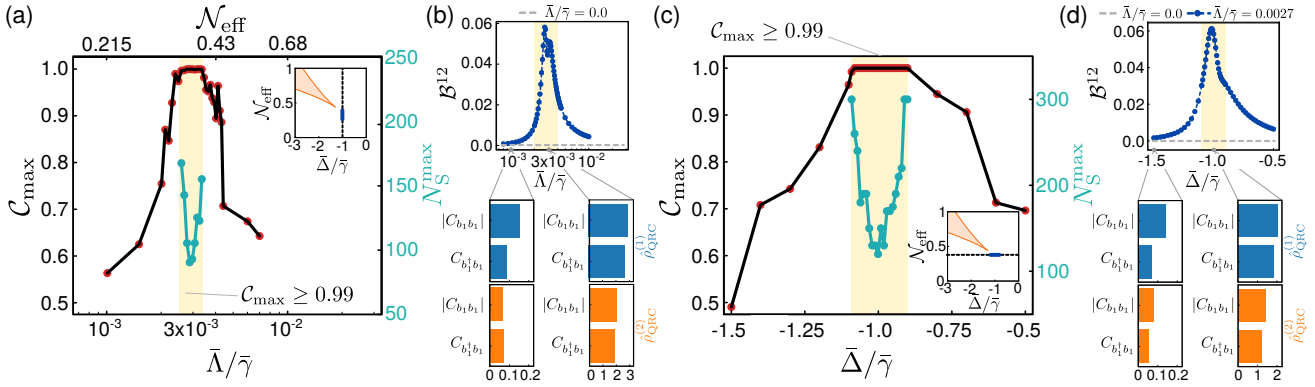


Figure 10. (a) C_{\max} (left-hand axis) as a function of nonlinearity $\bar{\Lambda}$, for fixed detuning $\bar{\Delta}/\bar{\gamma} = -1.0$. Where $C_{\max} \geq 0.99$ (yellow shaded region), N_S^{\max} is plotted along the right-hand axis (green). Inset indicates the classical phase diagram. (b) Unconditional QRC node amplitude difference $B^{12} = |\langle \hat{b}_1^{(1)} \rangle - \langle \hat{b}_1^{(2)} \rangle|$ calculated using TEOMs as a function of nonlinearity $\bar{\Lambda}$. Lower panel shows unconditional QRC cumulants corresponding to the states $\hat{\rho}_{\text{QRC}}^{(1),(2)}$ (blue and orange respectively) in a region of small B^{12} (left) and large B^{12} (right). Note the difference in magnitude of cumulants. (c), (d) Same as (a), (b) respectively, but now as a function of detuning $\bar{\Delta}$, for fixed nonlinearity $\bar{\Lambda}/\bar{\gamma} = 0.0027$, $\mathcal{N}_{\text{eff}} \simeq 0.354$.

form

$$\eta_{\text{eff}}^{(\sigma)} = -\frac{\Gamma_c}{\sqrt{2}} \left(\langle \hat{X}_{a_1}^{(\sigma)} \rangle + i \langle \hat{P}_{a_1}^{(\sigma)} \rangle \right) = \frac{2\Gamma_c \kappa \cdot \eta^{(\sigma)}}{4(G_{12}^{(\sigma)})^2 + \kappa(2G_1^{(\sigma)} - \kappa)} \quad (37)$$

where the second equality follows from the steady-state solutions for $\langle \hat{X}_{a_1}^{(\sigma)} \rangle$, $\langle \hat{P}_{a_1}^{(\sigma)} \rangle$ (see Appendix I). The current task is specifically chosen so that $\eta_{\text{eff}}^{(1)} = \eta_{\text{eff}}^{(2)}$; we hence set $\eta_{\text{eff}} \equiv |\eta_{\text{eff}}^{(1,2)}|$ to once again define an effective nonlinearity \mathcal{N}_{eff} that defines the classical phase diagram,

$$\mathcal{N}_{\text{eff}} = \frac{\eta_{\text{eff}}}{\bar{\gamma} + \Gamma_c} \sqrt{\frac{\bar{\Lambda}}{\bar{\gamma} + \Gamma_c}} \quad (38)$$

similar to Eq. (30), but now with $\bar{\gamma} \rightarrow \bar{\gamma} + \Gamma_c$ to account for the modified decay rate of the QRC due to the different form of coupling \mathcal{L}_c . The resulting phase diagram is plotted in the inset of Fig. 10(a). The dashed vertical line marks the cross-section along which Fig. 10(a) is calculated, while the solid blue line coincides with the yellow shaded region where $C_{\max} \geq 0.99$. Clearly, this region places the QRC in the vicinity of the classical bistability, where cumulants are amplified, precisely as observed in Sec. III for a single-node QRC under coherent driving.

2. Dependence on detuning $\bar{\Delta}$

Next we explore the classification performance as a function of detuning $\bar{\Delta}$ for fixed nonlinearity strength $\bar{\Lambda}/\bar{\gamma} = 0.0027$, which sets $\mathcal{N}_{\text{eff}} \simeq 0.354$ via Eq. (38). These results are shown in Fig. 10(c). Once more, we find that an optimal range of detuning values exists for which C_{\max} approaches unity and N_S^{\max} is minimized. Analysis of B^{12} as a function of detuning values in Fig. 10(d) shows how the QRC

node amplitude displacement peaks around a detuning value $\bar{\Delta}/\bar{\gamma} \simeq -1.0$ for the specific \mathcal{N}_{eff} value considered. This operating regime corresponds to $C_{\max} \geq 0.99$ (yellow shaded region), and once again places the QRC in the vicinity of the classical bistability (solid blue line, Fig. 10(c) inset).

3. Dependence on damping $\bar{\gamma}$

The QRC node damping rate appears as the natural scaling parameter for QRC hyperparameters, just as in the pointer state classification task. Hence if $\bar{\gamma}$ is varied, the remaining QRC hyperparameters $\{\mathcal{N}_{\text{eff}}, \bar{\Delta}/\bar{\gamma}\}$ can be suitably adjusted to ensure that the QRC operates in a nonlinear processing regime, enabling successful classification.

Beyond parameterizing such nonlinear processing regimes, $\bar{\gamma}$ also influences the QRC node response time and the amplitude of QRC output fields for fixed input signal strengths. Recall that our setup of the amplifier state classification task - using measurement records only beyond a time t_0 once transients have settled - effectively analyzes the QRC's processing of *steady-state* signals from the two-mode measured quantum system. Hence considerations of the QRC responding promptly to a dynamically-evolving signal as in the pointer state classification task do not apply directly. Regardless, the QRC damping rate $\bar{\gamma}$ can influence the time t_0 if it is *slower* than the amplifier response rate for the states under consideration; more precisely, in this case we would choose $t_0 \gg 1/\bar{\gamma}$. This in principle increases the total required measurement time. On the other hand, larger $\bar{\gamma}$ for fixed strength of incident drives will suppress the amplitude of QRC output, which is in general detrimental to classification. These considerations imply that $\bar{\gamma}$ should ideally be chosen to be on the order of the response rate of the amplifier (or the measured quantum system in general).

4. Summary

The results of this section identify a range of QRC hyperparameters $\{\mathcal{N}_{\text{eff}}, \bar{\Delta}/\bar{\gamma}\}$ where the QRC is able to successfully classify amplifier states with equal means, finding a strong connection to operating points in the vicinity of the classical bistability of the single coherently-driven QRC node, similar to the pointer state classification task. Importantly, the mechanism of classification is distinct: the input-dependent enhancement of second-order cumulants of the QRC state, whose nonlinearity leads to this difference manifesting via the QRC node amplitudes.

D. Classical vs. quantum regimes of operation

As with the task of pointer state classification, we now analyze QRC performance for amplifier state classification across decreasing input drives to the measurement chain. To this end, we consider different instances of this task, defined by the same two-mode amplifier interaction strengths $G_1^{(\sigma)}, G_{12}^{(\sigma)}$, but different values of drive input to the measurement chain, $(\eta^{(1)}/\bar{\gamma}, \eta^{(2)}/\bar{\gamma}) \in \{(2.0, 3.2), (5.0, 8.0), (10.8, 12.0)\}$. Via Eq. (38), these instances yield effective QRC drive strengths $\eta_{\text{eff}}/\bar{\gamma} \in \{20.0, 12.5, 5.0\}$, for which we calculate the classification accuracy metrics \mathcal{C}_{max} and N_S^{max} as a function of nonlinearity $\bar{\Lambda}$. We have fixed $\bar{\Delta}/\bar{\gamma} = 0.0$ and $\bar{\gamma} = \kappa$ as before.

We plot \mathcal{C}_{max} in the lower panel of Fig. 11 for the corresponding values of $\eta_{\text{eff}}/\bar{\gamma}$ (solid blue, purple, and peach respectively), with the yellow shaded region indicating nonlinearity strengths for which $\mathcal{C}_{\text{max}} \geq 0.99$. Similar to the pointer state classification task, we see that for lower η_{eff} , the nonlinearity strength required for successful classification increases. This is captured by the form of effective nonlinearity \mathcal{N}_{eff} ; colored crosses indicate QRCs with the same value of $\mathcal{N}_{\text{eff}} \simeq 0.354$ for the different instances of this classification task, which all succeed at classification but require stronger $\bar{\Lambda}$ for weaker η_{eff} .

Decreasing the input power therefore requires QRCs with stronger nonlinearity, providing a natural probe of the quantum regime of QRC operation. To analyze QRC performance across this transition, we also plot N_S^{max} in the lower panel of Fig. 11 (corresponding lighter colors, right-hand axis). Interestingly, we find that for decreasing $\bar{\Lambda}$ indicating the classical limit of QRC operation, the amplifier state classification task becomes more difficult, as the number of shots N_S^{max} required to obtain $\mathcal{C}_{\text{max}} \geq 0.99$ increases. Conversely, for operation in more quantum regimes, N_S^{max} decreases, so that the task has a reduced resource cost.

To understand this observation further, we use the TEOMs to calculate the difference in the unconditional QRC node amplitude for the two states, $\mathcal{B}^{12} = |\langle \hat{b}_1^{(1)} \rangle - \langle \hat{b}_1^{(2)} \rangle|$. The result is plotted in the top panel of Fig. 11 for each considered instance of the classification task. We clearly see that \mathcal{B}^{12} increases as we move towards QRCs operating deeper in the quantum regime (blue to purple to peach). For conditional




Figure 11. \mathcal{C}_{max} (solid curves) and N_S^{max} (lighter curves, right-hand axis) as a function of average nonlinearity hyperparameter $\bar{\Lambda}$ for instances of the amplifier state classification task with different effective drive amplitudes η_{eff} as indicated. Crosses indicate QRCs with fixed $\mathcal{N}_{\text{eff}} \simeq 0.354$ for the three considered task instances; for weaker $\eta_{\text{eff}}/\bar{\gamma}$, corresponding QRCs have stronger nonlinearity and operate in more quantum regimes (see text for more details). Top panel plots the unconditional QRC node amplitude difference $\mathcal{B}^{12} = |\langle \hat{b}_1^{(1)} \rangle - \langle \hat{b}_1^{(2)} \rangle|$, which increases for QRCs operating in more quantum regimes. The dashed gray line indicating the (vanishing) value of \mathcal{B}^{12} for linear QRCs.

evolution, this increase manifests in a larger relative displacement of measured quadrature distributions, resulting in a reduction in ensemble-averaging required for successful classification. Conversely, extrapolating the results to the classical limit ($\bar{\Lambda} \rightarrow 0$ for fixed \mathcal{N}_{eff}), we see that $\mathcal{B}^{12} \rightarrow 0$, so that the QRC is unable to distinguish the amplifier states in the classical limit.

We thus observe a qualitative difference in QRC performance across classical and quantum regimes when compared to the pointer state classification task. This is because unlike the latter, the information that distinguishes amplifier states is encoded in their second-order cumulants. In the classical limit governed primarily by first-order cumulants of the measurement chain, the influence of these second-order cumulants is suppressed; this directly reduces \mathcal{B}^{12} , rendering the classification task more difficult (requiring larger N_S^{max}). Conversely, in more quantum regimes the influence of second-order cumulants on QRC node amplitudes is enhanced, reducing N_S^{max} . These results indicate that the nature of input encoding into QRC evolution can play an important role in extracting possible processing advantages in quantum operating regimes.

However, one aspect of probing the quantum regime is shared across both classification tasks: the absolute amplitude of measured QRC outputs is reduced for weaker input drives in this regime. For the amplifier state classification task, the relative difference in QRC node amplitudes \mathcal{B}^{12} is still larger

in this regime, which is why the required N_S^{\max} decreases in our simulations of the complete measurement chain. In a real cQED experiment, however, the reduced amplitude can lead to the desired signal being swamped by noisy (classical) HEMT amplifiers that are also part of the measurement chain. As discussed earlier, this limitation is overcome by including quantum-limited amplifiers as part of the measurement chain, in this case to amplify the QRC output at cryogenic temperatures. We discuss the inclusion of such an amplifier for the current task in Appendix L. We find that even when using phase-preserving quantum-limited amplifiers that add unavoidable amplified quantum noise to the measured quadrature distributions, the associated amplification of the difference in QRC node amplitudes \mathcal{B}^{12} ensures that the benefit of operating the QRC in more quantum regimes is preserved.

VI. DISCUSSION AND OUTLOOK

In this paper we present and analyze a framework for quantum reservoir computing that we believe is timely within the current landscape of quantum computation. Research efforts geared towards the realization of a universal quantum computer to outperform classical computational paradigms have accelerated in recent years [74, 75]; however, several important challenges remain in scaling nascent quantum computers to larger and larger numbers of quantum degrees of freedom, in particular maintaining coherence times, correcting for the inevitable errors that occur during computation, and managing growing hardware complexity [76]. Nevertheless, as part of these and related efforts, modern experiments have continued to evolve in scope and complexity, realizing extended multi-mode nonlinear quantum systems with an unprecedented level of control. This naturally raises the question of whether such near-term quantum devices can already be used for nontrivial computation. This paper lays the theoretical foundations to assess the potential computational capabilities of physical quantum systems as they are currently realized, with constrained input and output control and subject to specific measurement schemes. The machine-learning paradigm of reservoir computing using physical systems, with relaxed requirements over internal device optimization, is ideally suited to harness these capabilities.

To this end, our strategy is to develop a framework that (1) is built upon a quantum-mechanical description of a physical quantum system as part of its complete quantum measurement chain, outputs from which are sampled via a prescribed measurement scheme, and (2) has a physically-motivated classical limit of operation. Such a framework allows us to determine performance metrics, such as fidelity and time-to-solution, of performing a computational task subject to operational constraints on inputs and outputs, and using only accessible information obtained via measurement, therefore accounting for experimental resource costs such as measurement time and repeated runs. Then, comparing these metrics across well-defined classical and quantum regimes of reservoir operation pinpoints advantages of quantum reservoir processing for a given computational task.

We believe such an approach can be applied broadly in seeking advantages or disadvantages to computation using quantum devices across very general computational tasks. For a concrete analysis, in this paper we realize this framework in the cQED architecture, and apply it to two quantum state classification tasks: states with distinct means (task I: pointer state classification), and states with equal means but distinct variances (task II: amplifier state classification). This requires developing an efficient computational approach to simulate the conditional dynamics of nonlinear quantum measurement chains incorporating Josephson-junctions, signal-generating quantum systems, with the entire chain evolving under continuous heterodyne measurement. We present such an approach based on a set of Stochastic Truncated Equations Of Motion (STEOMs) for measured observables and quantum cumulants. Different from earlier quantum reservoir computing approaches, this allows us to consider training of a linear output layer directly on measured outputs, which in addition to measurement noise include quantum fluctuations non-linearly transformed by the action of the quantum reservoir itself.

The imposition of a trainable *linear* output layer applied to obtained measurement records is an important choice: it allows any useful nonlinear processing that enables quantum state classification to be attributed to nonlinearity in the quantum measurement chain, and not to post-processing operations. In our realization, this nonlinearity is provided by the QRC, and is in fact critical: considering linear QRCs under the same operating conditions and output processing does not yield perfect classification. We note that the required nonlinearity could in principle be provided by other elements of a measurement chain, for example via alternate measurement schemes such as photodetection, or distinct input encoding schemes. This intuition is supported by work on classical reservoirs, which have demonstrated the power of computation with entirely linear reservoirs but employing non-linear readout layers [4, 77]. This flexibility speaks to the ability to implement physical reservoir computing across a range of quantum device platforms, tailored to specific available nonlinearities and measurement schemes. A comparison of the relative performance of different quantum reservoir computing implementations is an interesting avenue for future exploration.

The ability to analyze QRC performance across classical-to-quantum transitions within this framework provides a systematic approach to identifying the role of reservoir processing in regimes where the operation is governed by quantum statistical features. Carrying out this analysis, we find that pointer states, distinguished by their first-order cumulants, can be successfully classified even by QRCs operating in the classical limit. In addition, from the point of view of time-to-solution (t_{\max}), it is more advantageous to operate deeper in the classical limit. This may be perceived as a natural outcome of higher signal-to-noise ratio in the measured outputs, but the analysis shows the impact of non-linearly transformed noise should be considered carefully when reaching conclusions about the information processing capacity of a physical system in such limits. In contrast, the classification of amplifier states distinguished only by their second-order cumu-

lants are performed in a more sample-efficient manner (lower N_S^{\max}) deeper in the quantum regime of operation. From the point of view of power-efficiency, Task II provides an example where operating deeper in the quantum regime is advantageous but comes at the expense of the need for a QRC with a stronger non-linearity.

Finally, the computational efficiency of STEOMs permits us to train and test QRCs on sets of several stochastic measurement outputs across a wide range of parameter values beyond simply its bare nonlinearity. We identify regimes of operation where the QRC is successful at the chosen classification tasks, and find an interesting unifying feature: near classical bifurcation points, the QRC exhibits an enhanced non-linear response to its quantum inputs, which proves useful for quantum state classification. Such observations can aid in the design and operation of physical quantum reservoir computers.

Looking ahead, the presented framework is ideally suited to theoretical studies of the possible processing advantage afforded by quantum reservoirs when compared to their classical counterparts, across a variety of general computational tasks. This provides the basis for future studies of quantum reservoir processing capabilities as a function of reservoir size and hardware complexity. More importantly, the role of resources special to quantum reservoirs can be addressed: entanglement amongst the reservoir nodes, as well as between the quantum signal-generating system and the reservoir degrees of freedom. These studies are essential to identify the input and output schemes, and nature of quantum reservoirs that will be able to best extract the benefits of reservoir processing in the quantum regime. This is the subject of future work.

In view of practical implementations, two distinct information processing paradigms are made accessible: Firstly, it becomes possible to model quantum reservoirs as standalone quantum processors of data encoded in classical inputs, or quantum inputs via the state of a quantum system that is part of the same measurement chain and coupled to the quantum reservoir. Secondly, our analysis naturally provides a basis for studying quantum reservoirs as nonlinear quantum measurement devices, embedded into existing measurement chains for quantum systems, effectively describing an efficiently-trainable nonlinear readout layer. Our results indicate that simple QRCs can carry out useful nonlinear processing at cryogenic temperatures in a low noise environment, allowing experimentalists to avoid nonlinear processing of obtained noisy measurement records at room temperature.

ACKNOWLEDGMENTS

We would like to thank Dan Gauthier, Luke Govia, Michael Hatridge, Peter McMahon, Ioan Pop, Graham Rowlands and Guilhem Ribeill for useful discussions. This work is supported by AFOSR under Grant No. FA9550-20-1-0177 and the Army Research Office under Grant No. W911NF18-1-0144. Simulations in this paper were performed using the Princeton Research Computing resources at Princeton Univer-

sity, which is a consortium of groups led by the Princeton Institute for Computational Science and Engineering (PICSciE) and Office of Information Technology's Research Computing.

Appendix A: Quantum cumulants and the generating function

The starting point for our discussion of cumulants is via the characteristic function $\chi(z, z^*)$, defined in terms of the density operator $\hat{\rho}$ for an arbitrary quantum state as

$$\chi(z, z^*) = \text{tr} \left\{ \hat{\rho} e^{iz^* \hat{b}^\dagger} e^{iz \hat{b}} \right\} \quad (\text{A1})$$

The characteristic function as defined in Eq. (A1) has the property that it defines all *normal-ordered* operator averages with respect to the quantum state $\hat{\rho}$; simply evaluating the derivative of the characteristic function [54] at $z = z^* = 0$ yields

$$\frac{\partial^{p+q}}{\partial (iz^*)^p \partial (iz)^q} (\text{tr} \{ \hat{\rho} e^{iz^* \hat{b}^\dagger} e^{iz \hat{b}} \}) \Big|_{z=z^*=0} = \langle \hat{b}^{\dagger p} \hat{b}^q \rangle. \quad (\text{A2})$$

It is then possible to define a distribution function for calculating these normal-ordered moments,

$$P(\beta, \beta^\dagger) = \frac{1}{\pi^2} \int d^2 z \chi(z, z^*) e^{-iz^* \beta^*} e^{-iz \beta} \quad (\text{A3})$$

for which we use the suggestive notation $P(\zeta, \zeta^*)$, as it is easily shown to be none other than the Glauber-Sudarshan P distribution [54, 78, 79]:

$$\hat{\rho} = \int d^2 \beta P(\zeta, \zeta^\dagger) |\beta\rangle \langle \beta| \quad (\text{A4})$$

Having recalled the connection of the characteristic function to the quantum state defined by $\hat{\rho}$ and its normal-ordered moments, we now define a quantum generating function $G : \mathbb{C} \times \mathbb{C} \rightarrow \mathbb{R}$,

$$G(z, z^*) = \ln \text{tr} \{ \chi(z, z^*) \} = \ln \text{tr} \{ \hat{\rho} e^{iz^* \hat{b}^\dagger} e^{iz \hat{b}} \}. \quad (\text{A5})$$

The *quantum cumulants* are defined through Eq. (A5) by the series expansion of the generating function $G(z, z^*)$ which yields all quantum cumulants $C_{a^\dagger p a^q}$ for any $p, q \in \mathbb{N}^2$

$$\ln \text{tr} \{ \hat{\rho} e^{iz^* \hat{b}^\dagger} e^{iz \hat{b}} \} = \sum_{n=1}^{\infty} \sum_{\substack{p, q \in \mathbb{N} \\ p+q=n}} \frac{1}{p!q!} C_{b^{\dagger p} b^q} (iz^*)^p (iz)^q, \quad (\text{A6})$$

or equivalently,

$$C_{b^{\dagger p} b^q} \equiv \frac{\partial^{p+q} G}{\partial (iz^*)^p \partial (iz)^q} (z, z^*) \Big|_{z=z^*=0} \quad (\text{A7})$$

So far, cumulants may simply appear to be a mathematical construct. However, it can be shown that normal-ordered

cumulants and normal-ordered moments can be mapped to each other one-to-one. To see this, we can think of $G(z, z^*)$ as the composite of the logarithm function $\log(\chi) = \ln \chi$ and $\chi(z, z^*) = \text{tr} \left\{ \hat{\rho} e^{iz^* \hat{b}^\dagger} e^{iz \hat{b}} \right\}$, and employ multi-variable version of Faà di Bruno's formula for the n -th order partial derivative where $n = p + q$,

$$\begin{aligned} & \frac{\partial^{p+q} G}{\partial (iz^*)^p \partial (iz)^q} (z, z^*) \\ &= \sum_{\pi} \left(\log^{(|\pi|)}(\chi(z, z^*)) \prod_{B \in \pi} \frac{\partial^{p'+q'} \chi}{\partial (iz^*)^{p'} \partial (iz)^{q'}} (z, z^*) \right), \end{aligned} \quad (\text{A8})$$

where π defines the set of possible partitions of the ordered set $\{iz^*, \dots, iz^*, iz, \dots, iz\}$ with iz^* appearing p times and iz appearing q times. B indicates elements in this set of partitions, and p' and q' respectively counts the number of iz^* and iz in each partition element B . Finally, $f^{(|\pi|)}(\cdot)$ denotes the $|\pi|$ -order derivative of the function $f(\cdot)$, where $|\pi|$ denotes the cardinality of the set π . By noting that $\chi(0, 0) = \text{tr}(\hat{\rho}) = 1$, and then using $\log^{(|\pi|)}(\chi(0, 0)) = (\ln \chi)^{(|\pi|)}|_{\chi=1} = (-1)^{|\pi|-1} (|\pi| - 1)!$ we obtain

$$\begin{aligned} C_{b^{\dagger p} b^q} &= \frac{\partial^{p+q} G}{\partial (iz^*)^p \partial (iz)^q} (z, z^*) \Big|_{z=z^*=0} \\ &= \sum_{\pi} \left((-1)^{|\pi|-1} (|\pi| - 1)! \prod_{B \in \pi} \langle \hat{b}^{\dagger p'} \hat{b}^{q'} \rangle \right). \end{aligned} \quad (\text{A9})$$

As such, cumulants can be thought of as a reparameterization of the normal-ordered moments that typically appear in observables measured in an experiment. We will see in the next two appendices that this reparameterization leads to a much more efficient description of certain quantum states, which can then be leveraged to construct an efficient computational framework.

For our purposes, we will find it more useful to express quantum moments in terms of quantum cumulants. This inverse transformation can be performed by thinking of $\chi(z, z^*) = \text{tr} \left\{ \hat{\rho} e^{iz^* \hat{b}^\dagger} e^{iz \hat{b}} \right\}$ as the composite of the exponential function $\exp(G) = e^G$ and $G(z, z^*) = \ln \text{tr} \left\{ \hat{\rho} e^{iz^* \hat{b}^\dagger} e^{iz \hat{b}} \right\}$. Again, Faà di Bruno's formula gives

$$\frac{\partial^{p+q} \chi}{\partial (iz^*)^p \partial (iz)^q} (z, z^*) \quad (\text{A10})$$

$$= \sum_{\pi} \left(\exp^{(|\pi|)}(G(z, z^*)) \prod_{B \in \pi} \frac{\partial^{p'+q'} G}{\partial (iz^*)^{p'} \partial (iz)^{q'}} (z, z^*) \right) \quad (\text{A11})$$

where π defines the set of possible partitions of the ordered set $\{iz^*, \dots, iz^*, iz, \dots, iz\}$ with iz^* appearing p times and iz appearing q times, B indicates elements in this set of partitions, and p' and q' respectively counts the number of iz^* and iz in each partition element B . Again, noting that $G(0, 0) = \ln \text{tr}(\hat{\rho}) = 0$, and then using $\exp^{(|\pi|)}(G(0, 0)) =$

$(e^G)^{(|\pi|)}|_{G=0} = 1$ we get for n -th order partial derivative where $n = p + q$:

$$\begin{aligned} \langle \hat{b}^{\dagger p} \hat{b}^q \rangle &= \frac{\partial^{p+q} \chi}{\partial (iz^*)^p \partial (iz)^q} (z, z^*) \Big|_{z=z^*=0} \\ &= \sum_{\pi} \prod_{B \in \pi} \frac{\partial^{p'+q'} G}{\partial (iz^*)^{p'} \partial (iz)^{q'}} (z, z^*) \Big|_{z=z^*=0} \\ &= \sum_{\pi} \prod_{B \in \pi} C_{b^{\dagger p'} b^{q'}}. \end{aligned} \quad (\text{A12})$$

Finally, it is straightforward to generalize the definition of quantum cumulants to the N -node bosonic mode case. For a complex-valued vector $\mathbf{z} = (z_1, z_2, \dots, z_N)$, we can define the generating function $G: \mathbb{C}^N \times \mathbb{C}^N \rightarrow \mathbb{R}$ as:

$$G(\mathbf{z}, \mathbf{z}^*) = \ln \text{tr} \left\{ \hat{\rho} e^{iz_1^* \hat{b}_1^\dagger} \dots e^{iz_N^* \hat{b}_N^\dagger} e^{iz_1 \hat{b}_1} \dots e^{iz_N \hat{b}_N} \right\}. \quad (\text{A13})$$

The series expansion of $G(\mathbf{z}, \mathbf{z}^*)$ yields all quantum cumulants

$$\begin{aligned} & \frac{\partial^n G}{\partial (iz_1^*)^{p_1} \dots \partial (iz_N^*)^{p_N} \partial (iz_1)^{q_1} \dots \partial (iz_N)^{q_N}} (\mathbf{z}, \mathbf{z}^*) \Big|_{\mathbf{z}=\mathbf{z}^*=0} \\ &= C_{b_1^{\dagger p_1} \dots b_N^{\dagger p_N} b_1^{q_1} \dots b_N^{q_N}}. \end{aligned} \quad (\text{A14})$$

where $p_1 + \dots + p_N + q_1 + \dots + q_N = n$. The transformation between moments and cumulants in multi-node case can be obtained by simply replacing all single-node moments and cumulants in Eq.(A9), (A12) with multi-node moments and cumulants corresponding to normal-ordered operator $\hat{o}_1 \hat{o}_2 \dots \hat{o}_n = \hat{b}_1^{\dagger p_1} \dots \hat{b}_N^{\dagger p_N} \hat{b}_1^{q_1} \dots \hat{b}_N^{q_N}$:

$$C_{o_1 o_2 \dots o_n} = \sum_{\pi} \left((-1)^{|\pi|-1} (|\pi| - 1)! \prod_{B \in \pi} \langle \hat{o}_i : i \in B \rangle \right), \quad (\text{A15})$$

$$\langle \hat{o}_1 \hat{o}_2 \dots \hat{o}_n \rangle = \sum_{\pi} \prod_{B \in \pi} C_{o_i : i \in B}, \quad (\text{A16})$$

where π defines the set of possible partitions of operators in the n -order moment, and B indicates elements in this set of partitions.

Appendix B: Cumulants for some simple quantum states

Eq.(A9), (A12) and their multi-node case generalization Eq.(A15, A16) show that quantum cumulants and quantum moments have one-to-one correspondence through logarithm and exponential. On the other hand, for two complex-valued vectors $\mathbf{z} = (z_1, z_2, \dots, z_N)$ and $\boldsymbol{\beta} = (\beta_1, \beta_2, \dots, \beta_N)$, the P -representation of quantum state and quantum characteristic function are mutually related by Fourier transformation [54]:

$$P(\boldsymbol{\beta}, \boldsymbol{\beta}^*) = \frac{1}{\pi^2} \int \chi(\mathbf{z}, \mathbf{z}^*) e^{-iz \cdot \boldsymbol{\beta}^*} e^{-iz^* \cdot \boldsymbol{\beta}} d^2 \mathbf{z}, \quad (\text{B1})$$

$$\chi(\mathbf{z}, \mathbf{z}^*) = \int P(\boldsymbol{\beta}, \boldsymbol{\beta}^*) e^{iz \cdot \boldsymbol{\beta}^*} e^{iz^* \cdot \boldsymbol{\beta}} d^2 \boldsymbol{\beta}. \quad (\text{B2})$$

where the inner product is conventionally $\mathbf{z} \cdot \boldsymbol{\beta}^* = z_1^* \beta_1^* + \dots + z_N^* \beta_N^*$ and $\mathbf{z}^* \cdot \boldsymbol{\beta} = z_1 \beta_1 + \dots + z_N \beta_N$. These result ensures the one-to-one correspondence between density matrix and quantum cumulants.

Importantly, the truncation of cumulants to certain discrete orders naturally characterizes some special types of quantum states. In this section, we will show that states with only nonzero first order cumulants correspond to coherent states, while states with only nonzero first and second cumulants orders correspond to Gaussian states.

We first consider a product of coherent states $|\boldsymbol{\beta}_0\rangle \equiv |\beta_1, \beta_2, \dots, \beta_N\rangle$. In Eq. (A13), we then set $\hat{\rho} = |\boldsymbol{\beta}_0\rangle \langle \boldsymbol{\beta}_0|$, following which the RHS gives

$$\begin{aligned} & \ln \text{tr} \left\{ |\boldsymbol{\beta}_0\rangle \langle \boldsymbol{\beta}_0| e^{iz_1^* \hat{b}_1^\dagger} \dots e^{iz_N^* \hat{b}_N^\dagger} e^{iz_1 \hat{b}_1} \dots e^{iz_N \hat{b}_N} \right\} \\ &= \ln \left(\langle \boldsymbol{\beta}_0| e^{iz_1^* \hat{b}_1^\dagger} \dots e^{iz_N^* \hat{b}_N^\dagger} e^{iz_1 \hat{b}_1} \dots e^{iz_N \hat{b}_N} |\boldsymbol{\beta}_0\rangle \right) \\ &= \ln e^{iz_1^* \beta_1^* + \dots + iz_N^* \beta_N^* + iz_1 \beta_1 + \dots + iz_N \beta_N} \\ &= iz_1^* \beta_1^* + \dots + iz_N^* \beta_N^* + iz_1 \beta_1 + \dots + iz_N \beta_N, \quad (\text{B3}) \end{aligned}$$

From Eq. (A14), we obtain non-trivial derivative contributions only for first-order derivatives. Thus for all $k \in \{1, 2, \dots, N\}$, we have $C_{b_k} = \beta_k$ and $C_{b_k^\dagger} = \beta_k^*$, so that

$$(C_{b_1}, C_{b_2}, \dots, C_{b_N}) = \boldsymbol{\beta}_0, \quad (\text{B4a})$$

$$(C_{b_1^\dagger}, C_{b_2^\dagger}, \dots, C_{b_N^\dagger}) = \boldsymbol{\beta}_0^* \quad (\text{B4b})$$

We can also show that physical states with vanishing cumulants of order greater than one are necessarily coherent states. We begin with the normal-ordered characteristic function for such states, which is given by $\chi(\mathbf{z}, \mathbf{z}^*) = e^{G(i\mathbf{z}^*, i\mathbf{z})} = e^{i\mathbf{z} \cdot \boldsymbol{\beta}_0^* + i\mathbf{z}^* \cdot \boldsymbol{\beta}_0}$. The corresponding P -representation is simply the Fourier transform of the characteristic function, and is given by

$$\begin{aligned} P(\boldsymbol{\beta}, \boldsymbol{\beta}^*) &= \frac{1}{\pi^2} \int \chi(\mathbf{z}, \mathbf{z}^*) e^{-i\mathbf{z} \cdot \boldsymbol{\beta}^*} e^{-i\mathbf{z}^* \cdot \boldsymbol{\beta}} d^2 \mathbf{z} \\ &= \frac{1}{\pi^2} \int e^{i\mathbf{z} \cdot \boldsymbol{\beta}_0^* + i\mathbf{z}^* \cdot \boldsymbol{\beta}_0} e^{-i\mathbf{z} \cdot \boldsymbol{\beta}^*} e^{-i\mathbf{z}^* \cdot \boldsymbol{\beta}} d^2 \mathbf{z} \\ &= \delta(\boldsymbol{\beta} - \boldsymbol{\beta}_0), \quad (\text{B5}) \end{aligned}$$

which describes a product of coherent states. As a result, the states with all cumulants or order greater than one vanishing are exactly the collection of all coherent states.

Next, we consider quantum states whose quasi-probability distribution is Gaussian. Recall that the P , Q , and Wigner phase-space representations are related by the Weierstrass transformation, which always maps one Gaussian distribution to another Gaussian distribution [54]:

$$W(\boldsymbol{\beta}, \boldsymbol{\beta}^*) = \frac{2}{\pi} \int d^2 \tilde{\boldsymbol{\beta}} e^{-2|\boldsymbol{\beta} - \tilde{\boldsymbol{\beta}}|^2} P(\tilde{\boldsymbol{\beta}}, \tilde{\boldsymbol{\beta}}^*), \quad (\text{B6})$$

$$Q(\boldsymbol{\beta}, \boldsymbol{\beta}^*) = \frac{2}{\pi} \int d^2 \tilde{\boldsymbol{\beta}} e^{-2|\boldsymbol{\beta} - \tilde{\boldsymbol{\beta}}|^2} W(\tilde{\boldsymbol{\beta}}, \tilde{\boldsymbol{\beta}}^*). \quad (\text{B7})$$

The relation between the P -representation and characteristic function in Eqs. (B1), (B2) respectively is via Fourier transform, which also maps one Gaussian distribution to another Gaussian distribution.

Via this series of transformations, we therefore see that a quantum state with a Gaussian Wigner distribution will also have a Gaussian characteristic function, and hence must be entirely defined by cumulants of up to second-order (by definition of a Gaussian distribution).

Appendix C: Working with cumulants - TEOMs for single coherently-driven Kerr oscillator

Our approach requires the expression of moments of a multimode quantum system in terms of their associated cumulants. Moments and cumulants are related via their generating functions, which allows one to write a general prescription for writing an arbitrary normal-ordered joint system moment of order n in terms of a series of cumulants of order 1 up to n . This takes the specific form given by Eq. (A16),

$$\langle \hat{o}_1 \hat{o}_2 \dots \hat{o}_n \rangle = \sum_{\pi} \prod_{B \in \pi} C_{o_i: i \in B}, \quad (\text{C1})$$

where π defines the set of possible partitions of operators in the n -order moment, and B indicates elements in this set of partitions. The rather opaque form can be simplified by considering a specific case. We will use the above prescription to obtain the Eqs. (10)-(11a) in the main text for the single-node QRC. This minimal system already includes the key nonlinear element of the QRC, and therefore the most complex terms that need to be analyzed will arise already in these equations. We start by writing down the equation for $\langle \dot{\hat{b}} \rangle = \text{tr}\{\mathcal{L}\hat{\rho}\hat{b}\}$,

$$\langle \dot{\hat{b}} \rangle = \left(i\Delta - \frac{\tilde{\gamma}}{2} \right) \langle \hat{b} \rangle + i\Lambda \langle \hat{b}^\dagger \hat{b} \hat{b} \rangle - i\eta. \quad (\text{C2})$$

To rewrite the above equation in terms of cumulants, we must express the the third-order moment $\langle \hat{b}^\dagger \hat{b} \hat{b} \rangle$ in terms of cumulants. The possible partitions of the third-order moment are given by: $\pi \in \{(\hat{b}^\dagger \hat{b} \hat{b}), (\hat{b}^\dagger \hat{b}, \hat{b}) \times 2, (\hat{b} \hat{b}, \hat{b}^\dagger), (\hat{b}^\dagger, \hat{b}, \hat{b})\}$. Then, using Eq. (C1), we obtain

$$\langle \hat{b}^\dagger \hat{b} \hat{b} \rangle = C_{b^\dagger b b} + 2C_{b^\dagger b} \langle \hat{b} \rangle + C_{bb} \langle \hat{b}^\dagger \rangle + \langle \hat{b}^\dagger \rangle \langle \hat{b} \rangle^2. \quad (\text{C3})$$

Using the above, we can now rewrite Eq. (C2) in terms of cumulants as

$$\begin{aligned} \langle \dot{\hat{b}} \rangle &= \left(i\Delta - \frac{\tilde{\gamma}}{2} \right) \langle \hat{b} \rangle + i\Lambda \langle \hat{b}^\dagger \rangle \langle \hat{b} \rangle^2 - i\eta \\ &\quad + i\Lambda \left(C_{bb} \langle \hat{b}^\dagger \rangle + 2C_{b^\dagger b} \langle \hat{b} \rangle + C_{b^\dagger b b} \right) \end{aligned} \quad (\text{C4})$$

which is as given in Eq. (10).

To similarly obtain dynamical equations for cumulants $C_{b^\dagger b}, C_{bb}$, we write equations of motion for the second-order moments $\langle \hat{b}^\dagger \hat{b} \rangle$ and $\langle \hat{b} \hat{b} \rangle$ respectively. Starting with the former, we find:

$$\begin{aligned} \frac{d}{dt} \langle \hat{b}^\dagger \hat{b} \rangle &= -\tilde{\gamma} \langle \hat{b}^\dagger \hat{b} \rangle + i\eta \langle \hat{b} \rangle - i\eta \langle \hat{b}^\dagger \rangle \\ &= -\tilde{\gamma} C_{b^\dagger b} - \tilde{\gamma} \langle \hat{b}^\dagger \rangle \langle \hat{b} \rangle + i\eta \langle \hat{b} \rangle - i\eta \langle \hat{b}^\dagger \rangle \end{aligned} \quad (\text{C5})$$

where in the second line we rewrite the second-order moment in terms of cumulants. From this we can easily write $dC_{b^\dagger b} = d\langle \hat{b}^\dagger \hat{b} \rangle - (d\langle \hat{b}^\dagger \rangle)\langle \hat{b} \rangle - \langle d\langle \hat{b} \rangle \rangle \langle \hat{b}^\dagger \rangle$, which yields the dynamical equation for $C_{b^\dagger b}$:

$$\begin{aligned} \dot{C}_{b^\dagger b} &= -\tilde{\gamma}C_{b^\dagger b} - i\Lambda \left(C_{bb} \langle \hat{b}^\dagger \rangle^2 - C_{bb}^* \langle \hat{b} \rangle^2 \right) \\ &\quad - i\Lambda \left(C_{b^\dagger bb} \langle \hat{b}^\dagger \rangle - C_{b^\dagger bb}^* \langle \hat{b} \rangle \right) \end{aligned} \quad (\text{C6})$$

which is as written in Eq. (11a) of the main text.

Finally, we write the equation of motion for $\langle \hat{b}\hat{b} \rangle$,

$$\begin{aligned} \frac{d}{dt} \langle \hat{b}\hat{b} \rangle &= (2i\Delta - \tilde{\gamma}) \langle \hat{b}\hat{b} \rangle - 2i\eta \langle \hat{b} \rangle + i\Lambda \langle \hat{b}\hat{b} \rangle + i2\Lambda \langle \hat{b}^\dagger \hat{b}\hat{b}\hat{b} \rangle \\ &= (2i\Delta - \tilde{\gamma} + i\Lambda) (C_{bb} + \langle \hat{b} \rangle \langle \hat{b} \rangle) - 2i\eta \langle \hat{b} \rangle + i2\Lambda \langle \hat{b}^\dagger \hat{b}\hat{b}\hat{b} \rangle \end{aligned} \quad (\text{C7})$$

where we once again rewrite second-order moments in terms of cumulants in the second line. We now need to express the remaining fourth-order moment $\langle \hat{b}^\dagger \hat{b}\hat{b}\hat{b} \rangle$ in terms of cumulants. The possible partitions of this fourth-order moment are given by $\pi \in \{(\hat{b}^\dagger \hat{b}\hat{b}\hat{b}), (\hat{b}^\dagger \hat{b}\hat{b}, \hat{b}) \times 3, (\hat{b}\hat{b}\hat{b}, \hat{b}^\dagger), (\hat{b}^\dagger \hat{b}, \hat{b}\hat{b}) \times 3, (\hat{b}^\dagger \hat{b}, \hat{b}, \hat{b}) \times 3, (\hat{b}\hat{b}, \hat{b}^\dagger, \hat{b}) \times 3, (\hat{b}^\dagger, \hat{b}, \hat{b}, \hat{b})\}$, so that

$$\begin{aligned} \langle \hat{b}^\dagger \hat{b}\hat{b}\hat{b} \rangle &= C_{b^\dagger bbb} + 3\langle \hat{b} \rangle C_{b^\dagger bb} + \langle \hat{b}^\dagger \rangle C_{bbb} \\ &\quad + 3C_{b^\dagger b} C_{bb} + 3C_{b^\dagger b} \langle \hat{b} \rangle^2 + 3C_{bb} |\langle \hat{b} \rangle|^2 + \langle \hat{b}^\dagger \rangle \langle \hat{b} \rangle \langle \hat{b} \rangle \langle \hat{b} \rangle \end{aligned} \quad (\text{C8})$$

From this we can finally write $dC_{bb} = d\langle \hat{b}\hat{b} \rangle - 2(d\langle \hat{b} \rangle)\langle \hat{b} \rangle$, which yields the dynamical equation for C_{bb} :

$$\begin{aligned} \dot{C}_{bb} &= (2i\Delta - \tilde{\gamma} + i\Lambda) (C_{bb} + \langle \hat{b} \rangle^2) - 2i\eta \langle \hat{b} \rangle \\ &\quad + i2\Lambda \left(C_{b^\dagger bbb} + 3\langle \hat{b} \rangle C_{b^\dagger bb} + \langle \hat{b}^\dagger \rangle C_{bbb} \right) \\ &\quad + i2\Lambda \left(3C_{b^\dagger b} C_{bb} + 3C_{b^\dagger b} \langle \hat{b} \rangle^2 + 3C_{bb} |\langle \hat{b} \rangle|^2 + \langle \hat{b}^\dagger \rangle \langle \hat{b} \rangle^3 \right) \\ &\quad - 2 \left(i\Delta - \frac{\tilde{\gamma}}{2} \right) \langle \hat{b} \rangle^2 - i2\Lambda \langle \hat{b}^\dagger \rangle \langle \hat{b} \rangle^3 + 2i\eta \langle \hat{b} \rangle \\ &\quad - i2\Lambda \left(C_{bb} |\langle \hat{b} \rangle|^2 + 2C_{b^\dagger b} \langle \hat{b} \rangle^2 + C_{b^\dagger bb} \langle \hat{b} \rangle \right) \end{aligned} \quad (\text{C9})$$

which finally simplifies to:

$$\begin{aligned} \dot{C}_{bb} &= \\ & (2i\Delta - \tilde{\gamma} + i\Lambda) C_{bb} + i\Lambda \langle \hat{b} \rangle^2 + i4\Lambda |\langle \hat{b} \rangle|^2 C_{bb} + i6\Lambda C_{b^\dagger b} C_{bb} \\ & + i2\Lambda C_{b^\dagger b} \langle \hat{b} \rangle^2 + i2\Lambda C_{b^\dagger bbb} + i4\Lambda C_{b^\dagger bb} \langle \hat{b} \rangle + i2\Lambda C_{bbb} \langle \hat{b}^\dagger \rangle \end{aligned} \quad (\text{C10})$$

which can be re-arranged to the form written in Eq. (11b) of the main text.

Our computer algebra approach automates the above process of calculating contributions to equations of motion and expressing arbitrary moments in terms of cumulants, thus allowing the systematic truncation necessary to arrive at TEOMs (and STEOMs) introduced in the main text.

Appendix D: Calculating STEOMs - stochastic measurement contributions to truncated cumulant dynamics

To illustrate the salient features of the calculation, a single homodyne measurement superoperator suffices:

$$\mathcal{S}_{k,\text{meas}}(dW)\hat{\rho} = \sqrt{\frac{\gamma_k}{2}} \left(\hat{b}_k \hat{\rho} + \hat{\rho} \hat{b}_k^\dagger - \langle \hat{b}_k + \hat{b}_k^\dagger \rangle \right) dW_k^X(t) \quad (\text{D1})$$

For convenience, we will then write a general stochastic master equation in the form:

$$d\hat{\rho}^c = \mathcal{L}\hat{\rho}^c dt + \mathcal{S}_{k,\text{meas}}(dW)\hat{\rho}^c \quad (\text{D2})$$

where \mathcal{L} defines all contributions to the master equation governing deterministic evolution. We can thus write for the differential of conditional expectation of an arbitrary operator \hat{o} :

$$d\langle \hat{o} \rangle^c = \text{tr}\{\mathcal{L}\hat{\rho}^c \hat{o}\} dt + \sqrt{\frac{\gamma_k}{2}} \left(C_{ob_k}^c + C_{b_k^\dagger o}^c \right) dW_k^X(t) \quad (\text{D3})$$

Analogously, we can write for arbitrary second-order moments:

$$\begin{aligned} d\langle \hat{o}_1 \hat{o}_2 \rangle^c &= \text{tr}\{\mathcal{L}\hat{\rho}^c \hat{o}_1 \hat{o}_2\} dt + \sqrt{\frac{\gamma_k}{2}} \times \\ & \left(\langle \hat{o}_1 \hat{o}_2 \hat{b}_k \rangle^c + \langle \hat{b}_k^\dagger \hat{o}_1 \hat{o}_2 \rangle^c - \langle \hat{o}_1 \hat{o}_2 \rangle^c \langle \hat{b}_k \rangle^c - \langle \hat{b}_k^\dagger \rangle^c \langle \hat{o}_1 \hat{o}_2 \rangle^c \right) dW_k^X(t) \end{aligned} \quad (\text{D4})$$

We will now use the above results obtain the equation of motion for an arbitrary *normal-ordered* cumulant $C_{o_1 o_2}^c = \langle \hat{o}_1 \hat{o}_2 \rangle^c - \langle \hat{o}_1 \rangle^c \langle \hat{o}_2 \rangle^c$. The important step arises in writing the differential of this cumulant using Ito's lemma:

$$\begin{aligned} dC_{o_1 o_2}^c &= d\langle \hat{o}_1 \hat{o}_2 \rangle^c - (d\langle \hat{o}_1 \rangle^c) \langle \hat{o}_2 \rangle^c - \langle \hat{o}_1 \rangle^c (d\langle \hat{o}_2 \rangle^c) \\ &\quad - (d\langle \hat{o}_1 \rangle^c) (d\langle \hat{o}_2 \rangle^c) \end{aligned} \quad (\text{D5})$$

where the term on the second line arises from the modified chain rule in Ito calculus. We can now proceed to obtaining the individual contributions to the above equation.

We start with the first term in Eq. (D5), which was calculated in Eq. (D4). To proceed, it will prove convenient to write moments higher than first-order in terms of their corresponding cumulants. We begin with the second-order moments:

$$\begin{aligned} d\langle \hat{o}_1 \hat{o}_2 \rangle^c &= \text{tr}\{\mathcal{L}\hat{\rho}^c \hat{o}_1 \hat{o}_2\} dt + \\ & \sqrt{\frac{\gamma_k}{2}} \left(\langle \hat{o}_1 \hat{o}_2 \hat{b}_k \rangle^c + \langle \hat{b}_k^\dagger \hat{o}_1 \hat{o}_2 \rangle^c - C_{o_1 o_2}^c \langle \hat{b}_k \rangle^c - \langle \hat{b}_k^\dagger \rangle^c C_{o_1 o_2}^c \right. \\ & \left. - \langle \hat{o}_2 \rangle^c \langle \hat{o}_1 \rangle^c \langle \hat{b}_k \rangle^c - \langle \hat{o}_2 \rangle^c \langle \hat{o}_1 \rangle^c \langle \hat{b}_k^\dagger \rangle^c \right) dW_k^X(t) \end{aligned} \quad (\text{D6})$$

Next, we write down the expressions relating third-order moments to their corresponding cumulants using Eq. (C1):

$$\begin{aligned} C_{o_1 o_2 b_k}^c &= \langle \hat{o}_1 \hat{o}_2 \hat{b}_k \rangle^c - C_{o_1 o_2}^c \langle \hat{b}_k \rangle^c - C_{o_1 b_k}^c \langle \hat{o}_2 \rangle^c - C_{o_2 b_k}^c \langle \hat{o}_1 \rangle^c \\ &\quad - \langle \hat{o}_2 \rangle^c \langle \hat{o}_1 \rangle^c \langle \hat{b}_k \rangle^c \end{aligned} \quad (\text{D7a})$$

$$\begin{aligned} C_{b_k^\dagger o_1 o_2}^c &= \langle \hat{b}_k^\dagger \hat{o}_1 \hat{o}_2 \rangle^c - C_{o_1 o_2}^c \langle \hat{b}_k^\dagger \rangle^c - C_{b_k^\dagger o_1}^c \langle \hat{o}_2 \rangle^c - C_{b_k^\dagger o_2}^c \langle \hat{o}_1 \rangle^c \\ &\quad - \langle \hat{o}_2 \rangle^c \langle \hat{o}_1 \rangle^c \langle \hat{b}_k^\dagger \rangle^c \end{aligned} \quad (\text{D7b})$$

employing which, Eq. (D5) becomes:

$$\begin{aligned} d\langle\hat{o}_1\hat{o}_2\rangle^c &= \text{tr}\{\mathcal{L}\hat{\rho}^c\hat{o}_1\hat{o}_2\} dt + \\ &\sqrt{\frac{\gamma_k}{2}} \left(C_{o_1o_2b_k}^c + C_{o_1b_k}^c \langle\hat{o}_2\rangle^c + C_{o_2b_k}^c \langle\hat{o}_1\rangle^c \right) dW_k^X(t) \\ &+ \sqrt{\frac{\gamma_k}{2}} \left(C_{b_k^\dagger o_1o_2}^c + C_{b_k^\dagger o_1}^c \langle\hat{o}_2\rangle^c + C_{b_k^\dagger o_2}^c \langle\hat{o}_1\rangle^c \right) dW_k^X(t) \end{aligned} \quad (\text{D8})$$

where terms involving only first-order moments cancel. Note that the above includes terms of up to $O(dt)$, as dW_k^X is formally $O(dt^{1/2})$.

Next, we can calculate the second and third terms in Eq. (D5), using Eq. (D3). These take the simple forms:

$$\begin{aligned} (d\langle\hat{o}_1\rangle^c)\langle\hat{o}_2\rangle^c &= \langle\hat{o}_2\rangle^c \text{tr}\{\mathcal{L}\hat{\rho}^c\hat{o}_1\} dt \\ &+ \sqrt{\frac{\gamma_k}{2}} \left(C_{o_1b_k}^c \langle\hat{o}_2\rangle^c + C_{b_k^\dagger o_1}^c \langle\hat{o}_2\rangle^c \right) dW_k^X(t) \end{aligned} \quad (\text{D9a})$$

$$\begin{aligned} (d\langle\hat{o}_2\rangle^c)\langle\hat{o}_1\rangle^c &= \langle\hat{o}_1\rangle^c \text{tr}\{\mathcal{L}\hat{\rho}^c\hat{o}_2\} dt \\ &+ \sqrt{\frac{\gamma_k}{2}} \left(C_{o_2b_k}^c \langle\hat{o}_1\rangle^c + C_{b_k^\dagger o_2}^c \langle\hat{o}_1\rangle^c \right) dW_k^X(t) \end{aligned} \quad (\text{D9b})$$

Both equations above once again contain terms up to $O(dt)$.

Finally, we can write down the term arising from Ito's lemma. The requiring differential is that given by Eq. (D3). However, we need only retain the lowest $O(dt)$ term here, which is given by:

$$\begin{aligned} (d\langle\hat{o}_1\rangle^c)(d\langle\hat{o}_2\rangle^c) &= \\ \frac{\gamma_k}{2} \left(C_{o_1b_k}^c + C_{b_k^\dagger o_1}^c \right) \left(C_{o_2b_k}^c + C_{b_k^\dagger o_2}^c \right) dt + O(dt^{3/2}) \end{aligned} \quad (\text{D10})$$

where we have used $(dW_k^X)^2 = dt$.

Finally, we can write down Eq. (D5) by combining Eqs. (D8), (D9a), (D9b), and (D10). This finally yields:

$$\begin{aligned} dC_{o_1o_2}^c &= \text{tr}\{\mathcal{L}\hat{\rho}^c\hat{o}_1\hat{o}_2 - \langle\hat{o}_2\rangle^c \mathcal{L}\hat{\rho}^c\hat{o}_1 - \langle\hat{o}_1\rangle^c \mathcal{L}\hat{\rho}^c\hat{o}_2\} dt \\ &- \frac{\gamma_k}{2} \left(C_{o_1b_k}^c + C_{b_k^\dagger o_1}^c \right) \left(C_{o_2b_k}^c + C_{b_k^\dagger o_2}^c \right) dt \\ &+ \sqrt{\frac{\gamma_k}{2}} \left(C_{o_1o_2b_k}^c + C_{b_k^\dagger o_1o_2}^c \right) dW_k^X(t) \end{aligned} \quad (\text{D11})$$

The first term on the right hand side describes deterministic dynamics governed by \mathcal{L} . The second term is due to the measurement, but does not depend on the stochastic Wiener increment $dW_k^X(t)$. The last term is a stochastic contribution; however note that it is related only to third-order cumulants. Within the ansatz used in this paper, this explicitly stochastic term vanishes.

Appendix E: Complex- P representation of single coherently-driven QRC node

The complex- P representation associated phase-space variables (β, β^\dagger) with operators $(\hat{b}, \hat{b}^\dagger)$. For a single node of our

QRC model, defined as a single coherently-driven Kerr oscillator coupled to a zero temperature bath via Eq. (8), which we reproduce explicitly,

$$\begin{aligned} \mathcal{L}\hat{\rho} &= -i \left[-\Delta\hat{b}^\dagger\hat{b} - \frac{\Lambda}{2}\hat{b}^\dagger\hat{b}^\dagger\hat{b}\hat{b} + \eta(e^{-i\varphi_\eta}\hat{b} + e^{i\varphi_\eta}\hat{b}^\dagger), \hat{\rho} \right] \\ &+ (\gamma + \Gamma)\mathcal{D}[\hat{b}]\hat{\rho}. \end{aligned} \quad (\text{E1})$$

We note first the the drive phase φ_η can simply be absorbed into the definition of the operators $\hat{b} \rightarrow \hat{b}e^{i\varphi_\eta}$, $\hat{b}^\dagger \rightarrow \hat{b}^\dagger e^{-i\varphi_\eta}$, while leaving their commutator unchanged. The drive amplitude can thus be chosen to be completely real. The steady-state complex- P distribution $P_{\text{ss}}(\beta, \beta^\dagger)$ for the resulting system, in the phase space of variables β, β^\dagger associated with operators \hat{b}, \hat{b}^\dagger respectively, can be found exactly by the method of potentials [57, 80], and takes the form

$$\begin{aligned} P_{\text{ss}}(\beta, \beta^\dagger) &= \\ (\beta)^{(c-2)}(\beta^\dagger)^{(c^*-2)} \exp \left\{ \frac{2\eta}{\Lambda} \left(\frac{1}{\beta} + \frac{1}{\beta^\dagger} \right) + 2\beta^\dagger\beta \right\} \end{aligned} \quad (\text{E2})$$

where

$$c = \frac{-i\Delta + \frac{\gamma+\Gamma}{2}}{-i\Lambda}. \quad (\text{E3})$$

Knowledge of the exact steady-state complex P -distribution allows one to calculate arbitrary moments of the quantum steady-state of the driven nonlinear mode by integrating over complex phase space. We find for arbitrary normal-ordered steady-state moments [57]

$$\begin{aligned} \langle (\hat{b}^\dagger)^j (\hat{b})^i \rangle &= \\ e^{-i\varphi_\eta(i-j)} \left| \frac{2\eta}{\Lambda} \right|^2 \frac{\Gamma(c)\Gamma(c^*)}{\Gamma(c+i)\Gamma(c^*+j)} \frac{h(c+i, c^*+j, 8|\eta/\Lambda|^2)}{h(c, c^*, 8|\eta/\Lambda|^2)}, \end{aligned} \quad (\text{E4})$$

where we have reintroduced the drive phase by simply undoing the earlier transformation on operators \hat{b}, \hat{b}^\dagger . Here the function $h(x, y, z)$ is given by the hypergeometric series

$$h(x, y, z) = \sum_{n=0}^{\infty} \frac{z^n}{n!} \frac{\Gamma(x)\Gamma(y)}{\Gamma(x+n)\Gamma(y+n)} \quad (\text{E5})$$

and $\Gamma(x)$ is the Gamma function.

Eq. (E4) is used to calculate steady-state first-order moments and cumulants plotted in Fig. 3 in Sec. III of the main text.

Appendix F: Supplementary benchmarking simulations

In Fig. 4 of the main text, we provided results benchmarking the performance of STEOMs against full SME simulations for a single coherently-driven QRC node undergoing heterodyne measurement. In this appendix section, we present results for a more complex measurement chain that includes the

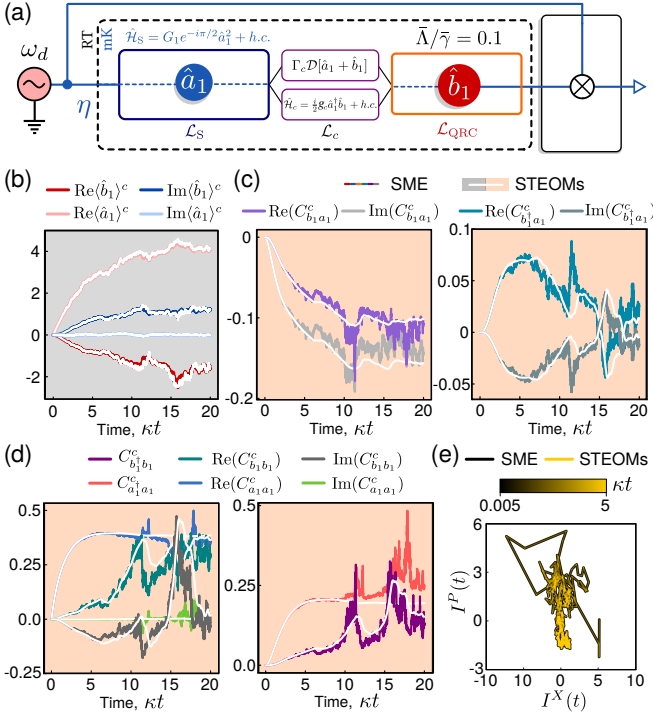


Figure 12. Benchmarking STEOMs against full SME simulations for the two-mode measurement chain depicted in (a). The single-mode measured system is an amplifier with parameters given by $\kappa_1 = \kappa/2$, $G_1/\kappa = 0.3$, and $\Delta_1/\kappa = 0.0$. The amplifier is coupled non-reciprocally (by setting $g_c = \Gamma_c = \kappa/2$) to a single-node QRC undergoing continuous measurement. The QRC parameters are given by $\bar{\gamma} = \kappa$, $\bar{\Lambda}/\bar{\gamma} = 0.1$, $\bar{\Delta}/\bar{\gamma} = -1.0$, such that $\mathcal{N}_{\text{eff}} = 0.385$. Simulation results for a single measurement trajectory showing (b) first-order cumulants, (c) second-order cross-cumulants, (d) second-order self-cumulants, and (e) measured quadratures $\{I^X(t), I^P(t)\}$.

key features common to the classification tasks considered in this paper: inclusion of a measured quantum system and its coupling to the QRC.

The specific measurement chain we consider is a simplified version of the full measurement chain employed for the task of amplifier state classification in Fig. 9(a) of the main text, and is depicted in Fig. 12(a). More precisely, it is defined by Eq. (1), with the system dynamics governed by Eq. (31) but with $G_{12} = 0.0$, thereby reducing the amplifier system to a single mode \hat{a}_1 . The rest of the measurement chain is unchanged: the mode \hat{a}_1 is coupled to the single QRC node via a directional hopping interaction given by Eq. (33), and the QRC node undergoes continuous measurement as defined by Eq. (2) with $K = 1$.

For fair benchmarking comparisons, we choose parameters where the QRC enables classification of the amplifier states; from results in the main text, we choose $\bar{\Delta}/\bar{\gamma} = -1.0$ and $\mathcal{N}_{\text{eff}} = 0.385$. Unfortunately, increasing the number of modes immediately makes the SME simulations much more computationally taxing. To lower the modal occupation numbers, we keep \mathcal{N}_{eff} fixed while increasing the nonlinearity strength to $\bar{\Lambda}/\bar{\gamma} = 0.1$, and simultaneously decreasing the drive strength to $\eta/\bar{\gamma} = 0.894$. We note that this nonlinearity strength is

about five times stronger than the largest nonlinearity used in the main text. All other parameters of the measurement chain are summarized in the caption of Fig. 12.

The complete measurement chain of $N_T = 2$ quantum modes is determined by $2N_T^2 + 3N_T = 14$ independent degrees of freedom: $\{\langle \hat{b}_1 \rangle^c, \langle \hat{a}_1 \rangle^c, C_{b_1 b_1}^c, C_{a_1 a_1}^c, C_{b_1 a_1}^c, C_{a_1 b_1}^c\}$ and their complex conjugates, as well as the real-valued cumulants $\{C_{b_1^\dagger b_1}^c, C_{a_1^\dagger a_1}^c\}$. We plot these quantities in Fig. 12(b)-(e) for a single measurement trajectory obtained using the STEOMs (white) and full SME simulations (color). We find that even though the nonlinearity is much stronger than those considered in the main text, there is in general good agreement between the methods. We also emphasize that for the same time step, SME simulations with a Hilbert space cut-off of 40 photons per quantum mode take about 37 hours to complete [62]; the STEOMs require 4 minutes.

Appendix G: Classical phase diagram for coherently-driven QRCs

In this appendix section we present the calculations required to obtain the classical phase diagram of the general two-node coherently-driven QRC analyzed in the main text. Our starting point is with the TEOMs describing the *unconditional* dynamics of the two-node QRC within the classical limit (i.e. also neglecting second-order cumulants):

$$\frac{d}{dt} \langle \hat{b}_1 \rangle = \left(+i\Delta_1 - \frac{\gamma_1}{2} \right) \langle \hat{b}_1 \rangle - ig_{12} \langle \hat{b}_2 \rangle - i\eta_1 + i\Lambda_1 \langle \hat{b}_1^\dagger \rangle \langle \hat{b}_1 \rangle, \quad (\text{G1})$$

$$\frac{d}{dt} \langle \hat{b}_2 \rangle = \left(+i\Delta_2 - \frac{\gamma_2}{2} \right) \langle \hat{b}_2 \rangle - ig_{12} \langle \hat{b}_1 \rangle - i\eta_2 + i\Lambda_2 \langle \hat{b}_2^\dagger \rangle \langle \hat{b}_2 \rangle. \quad (\text{G2})$$

We now introduce scaled variables that highlight the dependence of the classical dynamics of the two-node QRC on the nonlinearity strength and coherent drives. For convenience we set $\gamma_1 = \gamma_2 = \bar{\gamma}$, and once again perform the scaling introduced in Sec. III, introducing dimensionless time and analogous energy scales:

$$t' = \bar{\gamma}t, \quad (\text{G3a})$$

$$(\Delta'_k, \Lambda'_k, \eta'_k, g'_{12}) = (\Delta_k, \Lambda_k, \eta_k, g_{12})/\bar{\gamma} \quad (\text{G3b})$$

If we now scale $\langle \hat{b}_k \rangle \rightarrow \sqrt{\Lambda'_k} \langle \hat{b}_k \rangle'$, Eqs. (G1), (G2) transform to:

$$\frac{d}{dt'} \langle \hat{b}_1 \rangle' = \left(+i\Delta'_1 - \frac{1}{2} \right) \langle \hat{b}_1 \rangle' - ig'_{12} \sqrt{\frac{\Lambda'_1}{\Lambda'_2}} \langle \hat{b}_2 \rangle' - i\sqrt{\Lambda'_1} \eta'_1 + i\langle \hat{b}_1^\dagger \rangle' \langle \hat{b}_1 \rangle' \langle \hat{b}_1 \rangle', \quad (\text{G4})$$

$$\frac{d}{dt'} \langle \hat{b}_2 \rangle' = \left(+i\Delta'_2 - \frac{1}{2} \right) \langle \hat{b}_2 \rangle' - ig'_{12} \sqrt{\frac{\Lambda'_2}{\Lambda'_1}} \langle \hat{b}_1 \rangle' - i\sqrt{\Lambda'_2} \eta'_2 + i\langle \hat{b}_2^\dagger \rangle' \langle \hat{b}_2 \rangle' \langle \hat{b}_2 \rangle'. \quad (\text{G5})$$

It is immediately clear that the classical dynamics described by the above equations do *not* depend on the magnitude of

the nonlinearity strengths Λ'_k alone; instead, these appear as a scaling factor to the drive amplitudes, and via their *relative* strength Λ'_1/Λ'_2 . This indicates that *provided* Λ'_1/Λ'_2 is held fixed, decreasing Λ'_k while simultaneously increasing η'_k so that $\sqrt{\Lambda'_k}\eta'_k$ remain constant will leave the classical dynamics unchanged, apart from a scaling of node amplitudes $\langle \hat{b}_k \rangle'$. Analogously to the single-node QRC case analyzed in Sec. III of the main text, this transformation will reduce the influence of cumulants (not shown here), smoothly transitioning the two-node QRC to an effectively classical regime of operation.

We now specialize to the case of the two-node QRC analyzed in the main text, considering specific parameters, $\Delta'_1 = \Delta'_2 = \bar{\Delta}'$, $\Lambda'_1 = \Lambda'_2 = \bar{\Lambda}'$, and real coupling $g'_{12} = \bar{g}'$. Additionally, we consider only the \hat{b}_1 to be driven with a real-valued drive $(\eta'_1, \eta'_2) = (i\eta'_{\text{eff}}, 0)$, $\eta'_{\text{eff}} \in \mathbb{R}$, consistent with the coupling scheme employed in the main text. Then, the scaled TEOMs in Eqs. (G4), (G5) and their conjugates simplify to:

$$\begin{aligned} \frac{d}{dt'} \langle \hat{b}_1 \rangle' &= \left(+i\bar{\Delta}' - \frac{1}{2} \right) \langle \hat{b}_1 \rangle' - i\bar{g}' \langle \hat{b}_2 \rangle' + \mathcal{N}_{\text{eff}} \\ &\quad + i \langle \hat{b}_1^\dagger \rangle' \langle \hat{b}_1 \rangle' \langle \hat{b}_1 \rangle', \end{aligned} \quad (\text{G6})$$

$$\begin{aligned} \frac{d}{dt'} \langle \hat{b}_1^\dagger \rangle' &= \left(-i\bar{\Delta}' - \frac{1}{2} \right) \langle \hat{b}_1^\dagger \rangle' + i\bar{g}' \langle \hat{b}_2^\dagger \rangle' + \mathcal{N}_{\text{eff}} \\ &\quad - i \langle \hat{b}_1^\dagger \rangle' \langle \hat{b}_1^\dagger \rangle' \langle \hat{b}_1 \rangle', \end{aligned} \quad (\text{G7})$$

$$\begin{aligned} \frac{d}{dt'} \langle \hat{b}_2 \rangle' &= \left(+i\bar{\Delta}' - \frac{1}{2} \right) \langle \hat{b}_2 \rangle' - i\bar{g}' \langle \hat{b}_1 \rangle' \\ &\quad + i \langle \hat{b}_2^\dagger \rangle' \langle \hat{b}_2 \rangle' \langle \hat{b}_2 \rangle', \end{aligned} \quad (\text{G8})$$

$$\begin{aligned} \frac{d}{dt'} \langle \hat{b}_2^\dagger \rangle' &= \left(-i\bar{\Delta}' - \frac{1}{2} \right) \langle \hat{b}_2^\dagger \rangle' + i\bar{g}' \langle \hat{b}_1^\dagger \rangle' \\ &\quad - i \langle \hat{b}_2^\dagger \rangle' \langle \hat{b}_2^\dagger \rangle' \langle \hat{b}_2 \rangle'. \end{aligned} \quad (\text{G9})$$

where \mathcal{N}_{eff} takes the form given in Eq. (30) of the main text. We see once more that the classical dynamics of the two-node QRC depends on the drive and nonlinearity only via the effective nonlinear cooperativity \mathcal{N}_{eff} .

We now wish to analyze the stability of the two-node QRC and construct its classical phase diagram. To this end, we first obtain the steady state classical fixed points $(\langle \hat{b}_1 \rangle'_{\text{ss}}, \langle \hat{b}_2 \rangle'_{\text{ss}})$ of the two-node QRC by setting $\frac{d}{dt'} \langle \hat{b}_k \rangle'_{\text{ss}} = 0$ for $k = 1, 2$:

$$0 = \left(i\bar{\Delta}' - \frac{1}{2} \right) \langle \hat{b}_1 \rangle'_{\text{ss}} - i\bar{g}' \langle \hat{b}_2 \rangle'_{\text{ss}} + i |\langle \hat{b}_1 \rangle'_{\text{ss}}|^2 \langle \hat{b}_1 \rangle'_{\text{ss}} + \mathcal{N}_{\text{eff}}, \quad (\text{G10})$$

$$0 = \left(i\bar{\Delta}' - \frac{1}{2} \right) \langle \hat{b}_2 \rangle'_{\text{ss}} - i\bar{g}' \langle \hat{b}_1 \rangle'_{\text{ss}} + i |\langle \hat{b}_2 \rangle'_{\text{ss}}|^2 \langle \hat{b}_2 \rangle'_{\text{ss}}, \quad (\text{G11})$$

which describes a system of algebraic equations defining the steady-state amplitudes $\langle \hat{b}_k \rangle'_{\text{ss}}$ defining the classical fixed points.

To analyze the stability of these fixed points of the classical QRC dynamics, we evaluate the eigenvalues of Jacobian of the dynamical system. More explicitly, we treat the right hand side of Eq. (G6-G9) as a map $\mathbb{C}^4 \rightarrow \mathbb{C}^4$ with four variables $\langle \hat{b}_1 \rangle'_{\text{ss}}, \langle \hat{b}_1^\dagger \rangle'_{\text{ss}}, \langle \hat{b}_2 \rangle'_{\text{ss}}, \langle \hat{b}_2^\dagger \rangle'_{\text{ss}}$. The first-order derivative or Jacobian of this map gives

$$J[\langle \hat{b}_1 \rangle'_{\text{ss}}, \langle \hat{b}_2 \rangle'_{\text{ss}}] = \begin{pmatrix} -\frac{1}{2} + i\bar{\Delta}' + 2i|\langle \hat{b}_1 \rangle'_{\text{ss}}|^2 & i\langle \hat{b}_1 \rangle'^2_{\text{ss}} & -i\bar{g}' & 0 \\ -i\langle \hat{b}_1^\dagger \rangle'^2_{\text{ss}} & -\frac{1}{2} - i\bar{\Delta}' - 2i|\langle \hat{b}_1 \rangle'_{\text{ss}}|^2 & 0 & i\bar{g}' \\ -i\bar{g}' & 0 & -\frac{1}{2} + i\bar{\Delta}' + 2i|\langle \hat{b}_2 \rangle'_{\text{ss}}|^2 & i\langle \hat{b}_2 \rangle'^2_{\text{ss}} \\ 0 & i\bar{g}' & -i\langle \hat{b}_2^\dagger \rangle'^2_{\text{ss}} & -\frac{1}{2} - i\bar{\Delta}' - 2i|\langle \hat{b}_2 \rangle'_{\text{ss}}|^2 \end{pmatrix}. \quad (\text{G12})$$

Note that the Jacobian does not explicitly depend on the effective nonlinearity \mathcal{N}_{eff} ; this dependence enters implicitly via the steady state amplitudes $\langle \hat{b}_1 \rangle'_{\text{ss}}, \langle \hat{b}_2 \rangle'_{\text{ss}}$.

The matrix $J[\langle \hat{b}_1 \rangle'_{\text{ss}}, \langle \hat{b}_2 \rangle'_{\text{ss}}]$ has four complex eigenvalues $\lambda_s[\langle \hat{b}_1 \rangle'_{\text{ss}}, \langle \hat{b}_2 \rangle'_{\text{ss}}]$ where $s = 1, 2, 3, 4$ for each fixed point $\langle \hat{b}_1 \rangle'_{\text{ss}}, \langle \hat{b}_2 \rangle'_{\text{ss}}$. A fixed point is stable if and only if the largest real parts of all four eigenvalues is smaller than zero (indicat-

ing decaying fluctuations),

$$\max_{s=1,2,3,4} \text{Re}(\lambda_s[\langle \hat{b}_1 \rangle'_{\text{ss}}, \langle \hat{b}_2 \rangle'_{\text{ss}}]) < 0. \quad (\text{G13})$$

The classical phase diagram is thus simply constructed by calculating the total number of fixed points of the system for given QRC parameters and \mathcal{N}_{eff} , and then counting the number of stable fixed points. The classical bistable region corresponds to regions where two fixed points are stable. Note that

the single-node QRC phase diagram can be constructed using the above approach by setting $\bar{g} = 0$.

Appendix H: Pointer state classification as a function of QRC damping hyperparameter $\bar{\gamma}$

In this appendix section, we explore the role of the QRC nodes damping hyperparameter $\bar{\gamma}$ on the task of classification of cavity states studied in Sec. IV. We have already seen that the classification performance is strongly dependent on where the QRC is being operated with respect to its classical phase diagram, determined by parameters $\{\mathcal{N}_{\text{eff}}, \bar{g}/\bar{\gamma}, \bar{\Delta}/\bar{\gamma}\}$. The simplest role of $\bar{\gamma}$ is to simply modify these parameters and lead to operation of the QRC in a different part of the classical phase diagram, thereby affecting its performance in the same way as explored in Sec. IV.

However, we also know that $\bar{\gamma}$ controls the relaxation rate of QRC nodes and thus should play some role in the response of the QRC to time-dependent signals such as those relevant for the pointer state classification task. To be able to isolate this effect from that due to different operating points on the classical phase diagram, we must choose parameters that fix the QRC operating regime. This requires holding $\{\mathcal{N}_{\text{eff}}, \bar{g}/\bar{\gamma}, \bar{\Delta}/\bar{\gamma}\}$ all fixed while $\bar{\gamma}$ is varied. The coupling and detuning hyperparameters can be set independently of any other components in the measurement chain; we therefore set them according to

$$\frac{\bar{g}}{\bar{\gamma}} = 1.0, \quad \frac{\bar{\Delta}}{\bar{\gamma}} = 0.0, \quad (\text{H1})$$

while $\bar{\gamma}$ is varied (relative to κ). However, recall that the effective nonlinearity \mathcal{N}_{eff} , defined in Eq. (30),

$$\mathcal{N}_{\text{eff}} = \frac{\eta_{\text{eff}}}{\bar{\gamma}} \sqrt{\frac{\bar{\Lambda}}{\bar{\gamma}}} \quad (\text{H2})$$

depends on the cavity drive and detuning, and the amplification strength, via η_{eff} , in addition to QRC nonlinearity and damping hyperparameters. We must hold \mathcal{N}_{eff} fixed *without* modifying cavity or amplifier parameters (thus leaving η_{eff} unchanged), since changing either effectively changes the classification task required of the QRC, thus not allowing for a sensible comparison between QRC performance for different $\bar{\gamma}$ values. Then, when changing $\bar{\gamma}$, we simultaneously modify $\bar{\Lambda}$ such that \mathcal{N}_{eff} is held fixed. This is achieved by setting

$$\frac{\bar{\Lambda}}{\bar{\gamma}} = 0.005 \cdot \left(\frac{\bar{\gamma}}{\kappa}\right)^2 \implies \mathcal{N}_{\text{eff}} = \frac{\eta_{\text{eff}}}{\kappa} \sqrt{0.005} \simeq 1.27 \quad (\text{H3})$$

so that \mathcal{N}_{eff} is now constant regardless of the value of $\bar{\gamma}$; in calculating η_{eff} , we have used cavity parameters from Sec. IV. In effect, for a fixed drive η_{eff} , changing the damping rate modifies the QRC internal field amplitude, so that the QRC node nonlinearity must be increased to compensate and still achieve a strongly nonlinear response.

Of course the reduction in QRC field for increased damping rate will lead to a smaller QRC output field. This is true

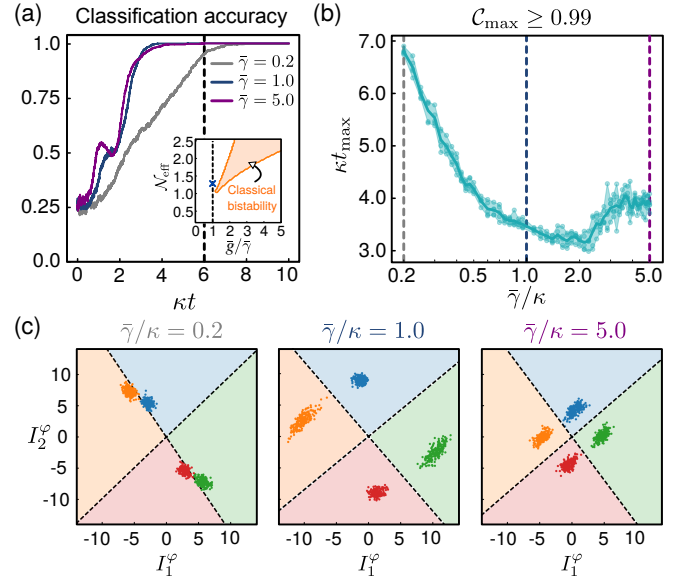


Figure 13. (a) Classification accuracy as a function of time for $\kappa t \in [0, 10]$ and specific damping hyperparameter values $\bar{\gamma} \in [0.2, 1.0, 5.0]\kappa$. Inset shows the classical phase diagram for the 2-node QRC, with the effective operating regime marked by the blue cross. (b) Time to classification κt_{max} as a function of damping hyperparameter $\bar{\gamma}$. (c) QRC outputs in measured phase space $\{I_1^\varphi, I_2^\varphi\}$ at time $\kappa t = 6.0$ for the specific damping hyperparameter values in (a).

even though in our formulation, the damping hyperparameter $\bar{\gamma}$ also controls the coupling to the output channel. As a result, the main consequence of our choice of scaling QRC hyperparameters is that we will remain on a fixed position of the *steady-state* classical phase diagram, but with increasing damping rate, the effective output field from the QRC will be reduced.

To analyze how these various considerations come in to play in determining QRC performance, we obtain the maximum classification accuracy C_{max} and the time to reach this maximum κt_{max} for the pointer state classification task as before. By virtue of fixing \mathcal{N}_{eff} so that we are in a nonlinear operating regime, we find that for all $\bar{\gamma}$ values considered, we reach $C_{\text{max}} \geq 0.99$ within the measurement time of $\kappa t \in [0, 10]$. Classification accuracy as a function of time at representative values of $\bar{\gamma} \in [0.2, 1.0, 5.0]\kappa$ is plotted in Fig. 13(a), indicating the approach to successful classification. From here, however, it is immediately clear that the performance of QRCs as a function of $\bar{\gamma}$ differ in the time taken to achieve perfect classification.

The extracted value of κt_{max} is plotted in Fig. 13(b) as a function of $\bar{\gamma}/\kappa$. We clearly observe a non-monotonic dependence of κt_{max} on the damping rate. For $\bar{\gamma} < \kappa$, t_{max} increases; this is expected, since the QRC is responding slower than the rate of evolution of the cavity signal. For larger $\bar{\gamma}$, the QRC response time is reduced, leading to a decrease in t_{max} . We clearly see that a minimum is reached around $\bar{\gamma} \simeq 2\kappa$. For larger $\bar{\gamma}$, we see that t_{max} begins to increase once more, even though the QRC response time is still fast relative to the

quantum system evolution timescale.

To understand this, we consider the measured phase space plot of integrated heterodyne records at time $\kappa t = 6.0$ plotted in Fig. 13(c) for the specific values of $\bar{\gamma} \in [0.2, 1.0, 5.0]\kappa$ from left to right. The plotted phase space area is the same across all panels. We clearly see that for all regimes of damping rate $\bar{\gamma}$ relative to κ , the QRC is operating in a nonlinear regime enabling classification of the four cavity signals, as arranged by fixing \mathcal{N}_{eff} . For $\bar{\gamma}/\kappa = 0.2$, the slow QRC response time means that this separation occurs much more slowly in time. For $\bar{\gamma}/\kappa = 5.0$, the states are well-separated into distinct areas of phase space; however we see clearly that the means of the output distributions are smaller than their values at lower QRC damping rates. This is because for fixed input signal strength η , the output field from the QRC is reduced in amplitude. As a result, the signal-to-noise ratio of measured quadratures is reduced for the same integration time. This manifests in a requirement of increased time t_{max} to successfully classify the cavity states.

Appendix I: Quantum states for amplifier state classification task: additional details

In this section, we will analyze the unconditional dynamics of the two-mode amplifier analyzed in Sec. V to justify parameter choices that general Gaussian states of the amplifier with equal first-order moments but distinct second-order cumulants.

We begin by writing down the *unconditional* equations of motion for the amplifier mode expectation values; these are given by

$$\begin{aligned} \frac{d}{dt}\langle\hat{a}_1\rangle &= \text{tr}\{\mathcal{L}_S\hat{\rho}\hat{a}_1\} + \text{tr}\{\mathcal{L}_c\hat{\rho}\hat{a}_1\} \\ &= \left(i\Delta_{a1} - \frac{\kappa_1 + \Gamma_c}{2}\right)\langle\hat{a}_1\rangle + G_1\langle\hat{a}_1^\dagger\rangle - iG_{12}\langle\hat{a}_2^\dagger\rangle + \eta \\ &\quad + \frac{g_c - \Gamma_c}{2}\langle\hat{b}_1\rangle, \end{aligned} \quad (\text{I1})$$

$$\begin{aligned} \frac{d}{dt}\langle\hat{a}_2\rangle &= \text{tr}\{\mathcal{L}_S\hat{\rho}\hat{a}_2\} + \text{tr}\{\mathcal{L}_c\hat{\rho}\hat{a}_2\} \\ &= \left(i\Delta_{a2} - \frac{\kappa_2}{2}\right)\langle\hat{a}_2\rangle - iG_{12}\langle\hat{a}_1^\dagger\rangle. \end{aligned} \quad (\text{I2})$$

Requiring a directional interaction between the two-mode amplifier and the QRC, we impose $g_c = \Gamma_c$. Furthermore, we now set $\kappa_1 + \Gamma_c = \kappa_2 \equiv \kappa$, so that the effective damping rate of both modes \hat{a}_1 and \hat{a}_2 is equal. Finally, for simplicity we also choose $\kappa_1 = \Gamma_c = \kappa/2$ and $\Delta_{a1} = \Delta_{a2} = 0$. Then, the simplified system above can be conveniently written in the quadrature basis as

$$\frac{d}{dt} \begin{pmatrix} \langle\hat{X}_{a1}\rangle \\ \langle\hat{P}_{a1}\rangle \\ \langle\hat{P}_{a2}\rangle \\ \langle\hat{X}_{a2}\rangle \end{pmatrix} = \mathbf{M} \begin{pmatrix} \langle\hat{X}_{a1}\rangle \\ \langle\hat{P}_{a1}\rangle \\ \langle\hat{X}_{a2}\rangle \\ \langle\hat{P}_{a2}\rangle \end{pmatrix} + \begin{pmatrix} \sqrt{2}\eta \\ 0 \\ 0 \\ 0 \end{pmatrix}, \quad (\text{I3})$$

where the dynamical matrix \mathbf{M} takes the form:

$$\mathbf{M} = \begin{pmatrix} G_1 - \frac{\kappa}{2} & 0 & 0 & -G_{12} \\ 0 & -G_1 - \frac{\kappa}{2} & -G_{12} & 0 \\ 0 & -G_{12} & -\frac{\kappa}{2} & 0 \\ -G_{12} & 0 & 0 & -\frac{\kappa}{2} \end{pmatrix}. \quad (\text{I4})$$

The form of \mathbf{M} already clarifies that the interaction $\propto G_1$ amplifies the \hat{X}_{a1} quadrature and squeezes the \hat{P}_{a1} quadrature. We are interested in the steady-state value of the quadrature expectation value $\langle\hat{X}_{a1}\rangle$ that couples to the QRC. This is easily found by setting the left hand side of Eq. (I3) to zero and solving the resulting system

$$\begin{pmatrix} \langle\hat{X}_{a1}\rangle \\ \langle\hat{P}_{a1}\rangle \\ \langle\hat{X}_{a2}\rangle \\ \langle\hat{P}_{a2}\rangle \end{pmatrix} = -\mathbf{M}^{-1} \begin{pmatrix} \sqrt{2}\eta \\ 0 \\ 0 \\ 0 \end{pmatrix}. \quad (\text{I5})$$

This yields the steady-state value of $\langle\hat{X}_{a1}^{(\sigma)}\rangle$ for parameters corresponding to the j th state generated by the two-mode amplifier

$$\langle\hat{X}_{a1}^{(\sigma)}\rangle = -\frac{2\kappa}{4(G_{12}^{(\sigma)})^2 + \kappa(2G_1^{(\sigma)} - \kappa)} \cdot \sqrt{2}\eta^{(\sigma)}. \quad (\text{I6})$$

It is easy to check that for the chosen parameter values corresponding to Fig. 9, namely $(\eta^{(\sigma)}, G_1^{(\sigma)}, G_{12}^{(\sigma)}) \in \{(5.0, 0.3, 0.0), (8.0, 0.0, 0.3)\}\kappa$, we have $\langle\hat{X}_{a1}^{(1)}\rangle = \langle\hat{X}_{a1}^{(2)}\rangle$. By extracting the steady-state solution for $\langle\hat{P}_{a1}\rangle$ from Eq. (I5), we similarly find $\langle\hat{P}_{a1}^{(1)}\rangle = \langle\hat{P}_{a1}^{(2)}\rangle = 0$.

Appendix J: Training and testing details

In this section, we provide details on the linear output layer for the processing of measurement records obtained from the QRC measurement chain, as applied to a C -state classification task.

QRC outputs. – We begin by defining the QRC output vector $\mathbf{x}(t)$ spanning the measured phase space, constructed using measured quadratures,

$$\mathbf{x}(t) = \begin{pmatrix} I_1^X(t) \\ I_1^P(t) \\ \vdots \\ I_K^X(t) \\ I_K^P(t) \end{pmatrix} = \begin{pmatrix} \langle\frac{1}{t}\int_0^t d\tau J_1^X(\tau)\rangle_{N_S} \\ \langle\frac{1}{t}\int_0^t d\tau J_1^P(\tau)\rangle_{N_S} \\ \vdots \\ \langle\frac{1}{t}\int_0^t d\tau J_K^X(\tau)\rangle_{N_S} \\ \langle\frac{1}{t}\int_0^t d\tau J_K^P(\tau)\rangle_{N_S} \end{pmatrix}. \quad (\text{J1})$$

For a K -node QRC under heterodyne measurement, $\mathbf{x}(t)$ as defined above is thus a $2K$ -dimensional vector at each t . Single-shot readout corresponds to setting $N_S = 1$.

Note that individual runs of the measurement chain for each of the C classes comprising the classification tasks yield distinct stochastic trajectories of measured quadratures as defined in Eq. (J1). Thus, we can index QRC output vectors

$\mathbf{x}_{(q)}^{(\sigma)}(t)$ by the corresponding class label $\sigma = 1, 2, \dots, C$, and with the index q labelling the stochastic trajectory.

Once the measured QRC output is obtained, an output layer acts on these outputs to ideally return the corresponding state label σ ; the result of the computation can therefore formally be written as

$$\sigma = f_N\{\mathbf{W}_O \mathbf{x}_{(q)}^{(\sigma)}(t) + \mathbf{b}\} \quad (\text{J2})$$

where \mathbf{W}_O is the matrix of trainable, time-independent output weights, \mathbf{b} is a vector of trainable biases.

In Eq. (J2), $f_N\{\cdot\}$ is a normalizing function that maps the vector of measured QRC outputs to a discrete, scalar state label $\sigma \in [1, \dots, C]$. This mapping is carried out via two operations. First, the QRC outputs $\mathbf{x}_{(q)}^{(\sigma)}(t)$ are mapped to an intermediate C -dimensional target vector $\mathbf{y}_{(q)}^{(\sigma)}(t)$ employing a ‘one-hot’ encoding (conventional for classification tasks): The i th element of this target vector $\mathbf{y}_{(q)}^{(\sigma)}(t)$ is given by:

$$[\mathbf{y}_{(q)}^{(\sigma)}(t)]_i = \begin{cases} 1 \forall t \text{ if } i = \sigma, \\ 0 \forall t \text{ otherwise.} \end{cases} \quad (\text{J3})$$

Note that elements of target vectors with one-hot encoding sum to unity. The output layer mapping QRC outputs to target vectors therefore also includes a normalization function,

$$\mathbf{y}_{(q)}^{(\sigma)}(t) = S[\mathbf{W}_O \mathbf{x}_{(q)}^{(\sigma)}(t) + \mathbf{b}] \equiv f_S[\mathbf{x}_{(q)}^{(\sigma)}(t)] \quad (\text{J4})$$

where $S[\cdot]$ is the normalized exponential (or ‘softmax’) function, defined by its action on an input vector \mathbf{v} :

$$S[\mathbf{v}]_i = \frac{e^{v_i}}{\sum_j e^{v_j}} \quad (\text{J5})$$

This function therefore has the property that $\sum_i S[\mathbf{v}]_i = 1$; in other words, it normalizes entries of a vector \mathbf{v} it acts upon such that the elements of the resulting vector add up to unity.

The reason to define target vectors with one-hot encoding is the property that the corresponding class label σ is simply given by

$$\sigma = \arg \max \{\mathbf{y}_{(q)}^{(\sigma)}(t)\} \quad (\text{J6})$$

where $\arg \max\{\mathbf{y}\}$ provides the index of the element of \mathbf{y} with largest magnitude. Therefore, combining Eqs. (J6), (J4) completes the definition of a trainable output layer mapping vector QRC outputs to discrete state labels,

$$\sigma = \arg \max \{S[\mathbf{W}_O \mathbf{x}_{(q)}^{(\sigma)}(t) + \mathbf{b}]\} \quad (\text{J7})$$

Comparing with Eq. (J2), we see that $f_N\{\cdot\} \equiv \arg \max \{S[\cdot]\}$, and is therefore a *nonlinear* normalizing operation on measured QRC outputs. However, note that it is fixed (not trained), and hence does not increase training complexity. More importantly, by analyzing linear QRCs followed by the output layer defined by Eq. (J7), we show in the main text that the nonlinearity of $f_N\{\cdot\}$ is insufficient to successfully perform the classification tasks considered in this paper.

Training.— The goal of training can now be defined: to learn trainable weights \mathbf{W}_O and biases \mathbf{b} to optimally map measured QRC outputs $\mathbf{x}_{(q)}^{(\sigma)}(t)$ to target vectors $\mathbf{y}_{(q)}^{(\sigma)}(t)$; the fidelity of the mapping is determined by minimizing a suitable cost function. We now discuss how measured QRC outputs are assembled in a form that aids this optimization on a digital computer.

To this end, we first note that the time label in the above definitions is actually discrete; the measurement window $t \in [t_0, t_f]$ is sampled at M_T temporal grid points (determined by the temporal resolution of measurement apparatus in experiments, and by the temporal step size of integration of STEOMs in simulations). Any segment of this measurement window, defined by times $t \in [t_s, t_e]$ with a total of m_T grid points, can be used for training; for example, for the pointer state classification task in the main text, the entire measurement window was used (so $t_s = t_0, t_e = t_f$, and $m_T = M_T$), while for amplifier state measurement, only the final integrated time point was used ($t_s = t_e = t_f$, and $m_T = 1$). The specific choice can be task-dependent.

Thus we can define a matrix $\mathbf{X}_{(q)}$ for the q th run of the measurement chain for state σ :

$$\mathbf{X}_{(q)} = \begin{pmatrix} \mathbf{x}_{(q)}^{(\sigma=1)}(t_s) \cdots \mathbf{x}_{(q)}^{(\sigma=1)}(t_e) & \cdots & \mathbf{x}_{(q)}^{(\sigma=C)}(t_s) \cdots \mathbf{x}_{(q)}^{(\sigma=C)}(t_e) \end{pmatrix} \quad (\text{J8})$$

which is evidently a $2K$ -by- $(C \times m_T)$ matrix. The corresponding matrix of output vectors takes the analogous form:

$$\mathbf{Y}_{(q)} = \begin{pmatrix} \mathbf{y}_{(q)}^{(\sigma=1)}(t_s) \cdots \mathbf{y}_{(q)}^{(\sigma=1)}(t_e) & \cdots & \mathbf{y}_{(q)}^{(\sigma=C)}(t_s) \cdots \mathbf{y}_{(q)}^{(\sigma=C)}(t_e) \end{pmatrix} \quad (\text{J9})$$

and is a C -by- $(C \times m_T)$ matrix.

A training set is constructed by obtaining a total of Q_{Train} measured QRC outputs, indexed as $q = 1, \dots, Q_{\text{Train}}$. For measured quadratures constructed by averaging over N_S individual shots from the measurement chain, this requires $Q_{\text{Train}} \times N_S$ runs of the measurement chain per state σ . For each q we obtain a matrix of QRC outputs $\mathbf{X}_{(q)}$ and corresponding targets $\mathbf{Y}_{(q)}$. These can be compiled into a composite QRC output matrix \mathbf{X} by concatenation of the individual matrices $\mathbf{X}_{(q)}$:

$$\mathbf{X} = (\mathbf{X}_{(q=1)} \cdots \mathbf{X}_{(q=Q_{\text{Train}})}) \quad (\text{J10})$$

which is clearly a $2K$ -by- $(Q_{\text{Train}} \times C \times m_T)$ matrix. The corresponding target matrix \mathbf{Y} is similarly defined as:

$$\mathbf{Y} = (\mathbf{Y}_{(q=1)} \cdots \mathbf{Y}_{(q=Q_{\text{Train}})}) \quad (\text{J11})$$

which is a C -by- $(Q_{\text{Train}} \times C \times m_T)$ matrix.

The optimal matrix of weights $\mathbf{W}_O^{\text{opt}}$ and biases \mathbf{b}^{opt} are finally determined by least squares optimization:

$$\{\mathbf{W}_O^{\text{opt}}, \mathbf{b}^{\text{opt}}\} = \min_{\mathbf{W}_O, \mathbf{b}} \{\|\mathbf{Y} - f_S[\mathbf{X}]\|^2\} \quad (\text{J12})$$

where $f_S[\mathbf{X}]$ is to be interpreted as operating column-wise on \mathbf{X} , in accordance with Eq. (J4).

A final note before discussing the results of this training procedure relates to the particular form of the measured subspace for the pointer state classification task analyzed in Sec. IV. The full measured phase space is projected onto a set of optimal quadratures determined by phases $\{\varphi_k\}$. This projection is achieved by simply rewriting the matrix of trainable weights \mathbf{W}_O as:

$$\mathbf{W}_O = \widetilde{\mathbf{W}}_O \mathcal{M}^\varphi \quad (\text{J13})$$

where \mathcal{M}^φ is the K -by- $2K$ matrix defined in terms of K phase angles introduced in Eq. (27) of the main text, and $\widetilde{\mathbf{W}}_O$ is now a C -by- K matrix of weights. With this reparameterization of \mathbf{W}_O , the rest of the training procedure remains unchanged, and the optimal matrix $\mathbf{W}_O^{\text{opt}}$ is still given by Eq. (J12). Formally, the only change is that the number of independent trainable weights in \mathbf{W}_O is reduced from $C \times 2K$ to $C \times K$ weights $\widetilde{\mathbf{W}}_O$ and K projection phases $\{\varphi_k\}$, and that the projection phases must be trained with bounds $[0, 2\pi]$.

Linear decision boundaries. – Once the weight matrix (and biases) are trained via Eq. (J12), they define linear decision boundaries that separate different classes and allow classification, as we will now show explicitly. We begin with the following observation: on the decision boundary that separates an arbitrary class i from class j , QRC outputs corresponding to $\sigma = i$ could be equally classified as belonging to $\sigma = j$, and vice versa. Mathematically, this implies $\mathbf{y}^{(\sigma=i)}(t) = \mathbf{y}^{(\sigma=j)}(t)$ on the decision boundary; from Eq. (J4), this condition can be written as:

$$S[\mathbf{W}_O \mathbf{x}(t) + \mathbf{b}]_i = S[\mathbf{W}_O \mathbf{x}(t) + \mathbf{b}]_j \quad (\text{J14})$$

or, using the definition of the softmax function in Eq. (J5),

$$\exp\{[\mathbf{W}_O \mathbf{x}(t) + \mathbf{b}]_i\} = \exp\{[\mathbf{W}_O \mathbf{x}(t) + \mathbf{b}]_j\} \quad (\text{J15})$$

where the normalization factor is common to both sides and cancels. This immediately requires the arguments of the exponentials to be equal, a requirement that can be rewritten in the form:

$$(b_i - b_j) + \sum_{k=1}^{2K} ((\mathbf{W}_O)_{ik} - (\mathbf{W}_O)_{jk}) x_k = 0 \quad (\text{J16})$$

which is simply the definition of a hyperplane in the measured phase space spanned by $\{x_1, \dots, x_k, \dots, x_{2K}\}$. This hyperplane defines the linear decision boundary that separates arbitrary classes i and j .

We now specialize the above general prescription to the case of pointer state classification considered in the main text. First, we find that the bias terms vanish due to the symmetry of the task being considered. The minor change is the introduction of trainable phase angles $\{\varphi_1, \varphi_2\}$, as mentioned earlier. To proceed, we simply rewrite \mathbf{W}_O in Eq. (J16), in




Figure 14. (a) Convergence of trained optimal weights $\mathbf{W}_O^{\text{opt}}$ for the pointer state classification task as a function of training set size Q_{Train} , parameterized as defined in Eq. (J13) in terms of elements of $\widetilde{\mathbf{W}}_O^{\text{opt}}$ (left panel) and phases $\{\varphi_k^{\text{opt}}\}$ (right panel). (b) Convergence of κt_{max} using optimal trained weights as a function of training set size Q_{Train} .

accordance with Eq. (J13);

$$\begin{aligned} \sum_{l=1}^K \sum_{k=1}^{2K} ((\widetilde{\mathbf{W}}_O)_{il} \mathcal{M}_{lk}^\varphi - (\widetilde{\mathbf{W}}_O)_{jl} \mathcal{M}_{lk}^\varphi) x_k &= 0 \\ \implies \sum_{l=1}^K ((\widetilde{\mathbf{W}}_O)_{il} - (\widetilde{\mathbf{W}}_O)_{jl}) I_l^\varphi &= 0 \end{aligned} \quad (\text{J17})$$

In going from the first to the second line the sum over k is performed, and simply yields $\sum_k \mathcal{M}_{lk}^\varphi x_k = I_l^\varphi$ as defined by Eq. (27) of the main text. Recalling that $K = 2$, the linear decision boundaries separating classes i and j are now lines in the measured phase space $\{I_1^\varphi, I_2^\varphi\}$, given by:

$$I_2^\varphi = - \frac{(\widetilde{\mathbf{W}}_O)_{i1} - (\widetilde{\mathbf{W}}_O)_{j1}}{(\widetilde{\mathbf{W}}_O)_{i2} - (\widetilde{\mathbf{W}}_O)_{j2}} \cdot I_1^\varphi \quad (\text{J18})$$

Testing. – Once $\mathbf{W}_O^{\text{opt}}$, \mathbf{b}^{opt} have been learned, the now-trained output layer can be used to test classification performance. The class label $\sigma_{(q)}^{p(\sigma)}$ predicted from a measured QRC output trajectory $\{\mathbf{x}_{(q)}^{(\sigma)}\}$ that is part of the test set ($q = 1, \dots, Q_{\text{Test}}$) is obtained by passing the QRC trajectory through the trained output layer defined in Eq. (J7),

$$\sigma_{(q)}^{p(\sigma)} = \arg \max \{\mathbf{y}_{(q)}^{(\sigma)}\} = \arg \max \{S[\mathbf{W}_O^{\text{opt}} \mathbf{x}_{(q)}^{(\sigma)} + \mathbf{b}^{\text{opt}}]\}. \quad (\text{J19})$$

Classification is correct if $\sigma_{(q)}^{p(\sigma)} = \sigma$. We have suppressed time labels on measured QRC trajectories and predicted class labels for brevity.

By predicting class labels using the entire test set of size Q_{Test} , the mean classification accuracy can be computed as

$$\text{Classification accuracy} = \frac{1}{C} \sum_{\sigma=1}^C \left[\frac{1}{Q_{\text{Test}}} \sum_{q=1}^{Q_{\text{Test}}} \delta_{\sigma_{(q)}^{p(\sigma)}, \sigma} \right], \quad (\text{J20})$$

where $\delta_{i,j}$ is the Kronecker delta function, here used simply to count instances where $\sigma_{(q)}^{p(\sigma)} = \sigma$ indicating correct classification.

Convergence with training set size.— Having discussed the formal training protocol we analyze briefly how training results vary with training set size, applied to the pointer state classification task. The left panel of Fig. 14(a) shows elements of the trained C -by- K matrix $\widetilde{\mathbf{W}}_O^{\text{opt}}$ (solid lines are first column elements, dashed lines are second column elements) as a function of training set size Q_{Train} . The right panel shows trained projection phases $\{\varphi_k^{\text{opt}}\}$. In Fig. 14(b) we plot the metric t_{max} that characterizes classification performance on the *training* set (for the QRC used to obtain this training set, $\mathcal{C}_{\text{max}} \geq 0.99$). We see that the obtained optimal weights in particular change rapidly for short training set sizes Q_{Train} , with a corresponding irregular change in t_{max} . However with increasing training set size, both the the trained weights and, more importantly, t_{max} converge. We find that the optimization landscape for projection phases is relatively shallow, so that small changes in $\{\varphi_k^{\text{opt}}\}$ do not significantly change t_{max} . We use this approach of testing classification performance using the training data set to determine Q_{Train} required for training in our classification tasks.

Appendix K: Training on finitely-sampled measured distributions vs. training using expectation values

In this appendix section, we analyze the difference in training using expectation values and using noisy data sets of consisting of measured quadrature records. For simplicity, we consider the amplifier state classification task of Sec. V, for $\bar{\Lambda}/\bar{\gamma} = 0.02$. All other QRC and amplifier hyperparameters are as mentioned in Sec. V A.

We begin by recalling the approach to training in our quantum reservoir computing framework. Consider the measured quadrature distributions shown in measured phase space in (1), the left panel of Fig. 15. Each colored point is one of 1500 measured quadrature records per amplifier state (in orange and blue respectively), each constructed using $N_S = 10$ individual measurement records. The linear decision boundary learned based on the measured *distributions* is shown in dashed black. For (2), $N_S = 150$, the measured quadrature distributions have a reduced variance due to increased ensemble averaging. The decision boundary trained on these measured distributions is different from (1), being tailored to the experimental resources used in output processing. We can test the classification accuracy of trained QRCs as in Sec. V, using $Q_{\text{Test}} = 3000$ measured quadrature records per state, each averaged over N_S shots. This result is plotted in the right panel of Fig. 15 in solid black, and shows the same qualitative form familiar from Fig. 9.

We now consider how training on expectation values might be carried out. In practice, expectation values must also be obtained from measurement records. Here, we consider expectation values of the measured quadratures as being approximately obtained once a large number of measurement records are averaged over; here we consider $N_S = 3000$. In measured phase space, these expectation values are marked by the red and blue dots respectively in the left panel of Fig. 15. The linear decision boundary in this case is simply the perpendicular bisector of the line joining these two mean values, and is shown by the dashed gray line.

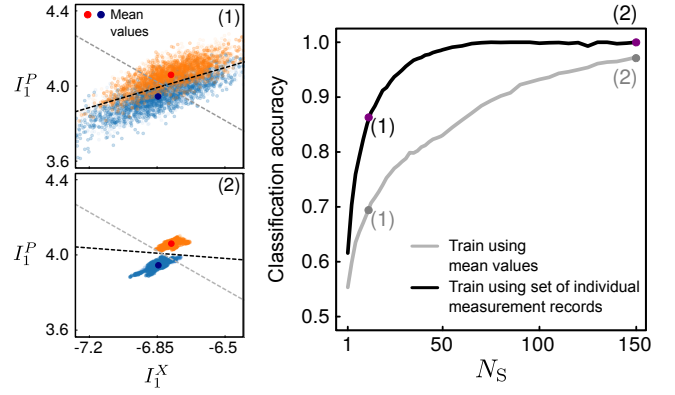


Figure 15. Classification accuracy as a function of number of shots N_S . Linear decision boundaries are obtained by training using either Q_{Train} measured quadrature records (black), or using only indicated mean values (gray).

ular bisector of the line joining these two mean values, and is shown by the dashed gray line.

The classification accuracy plotted in solid gray in the right panel of Fig. 15 corresponds to testing using this linear decision boundary learned using the two expectation values; we see that it underperforms relative to the case of linear decision boundaries trained using measured quadrature distributions. Note that the linear decision boundary learned using expectation values has no information about the non-isotropic (i.e. squeezed) distribution of measured quadratures; how this is detrimental is obvious from the distributions and decision boundaries in (1) and (2). Note that gray curve asymptotically approaches the black curve as $N_S \rightarrow \infty$ limit. In this limit, the measured distributions approach their expectation values, and the difference in the two training protocols is suppressed.

Appendix L: Post-Amplification

In this appendix section we consider the amplifier state classification task discussed in Sec. V of the main text, but now including a post-amplifier (PA), directionally-coupled to the single-node QRC, to amplify the QRC output. The corresponding schematic of the complete measurement chain is shown in Fig. 16(a).

We model the PA as a non-degenerate (phase-preserving) quantum-limited amplifier consisting of two modes \hat{d}_1, \hat{d}_2 . The inclusion of the PA modifies the Liouvillian describing the QRC. The SME describing the measurement chain is now given by

$$d\hat{\rho}^c = \mathcal{L}_S \hat{\rho}^c dt + \mathcal{L}_c \hat{\rho}^c dt + \mathcal{L}_{\text{QRC-PA}} \hat{\rho}^c dt + \widetilde{\mathcal{S}}_{\text{meas}}(dW) \hat{\rho}^c. \quad (\text{L1})$$

\mathcal{L}_S and \mathcal{L}_c are as defined by Eq. (31) and Eq. (33) of the main text respectively. The modified Liouvillian $\mathcal{L}_{\text{QRC-PA}}$ describes the evolution of the single-node QRC coupled to the PA, and is given by

$$\mathcal{L}_{\text{QRC-PA}} \hat{\rho} = \widetilde{\mathcal{L}}_{\text{QRC}} \hat{\rho} + \widetilde{\mathcal{L}}_c \hat{\rho} + \mathcal{L}_{\text{PA}} \hat{\rho}. \quad (\text{L2})$$

Here, $\tilde{\mathcal{L}}_{\text{QRC}}$ describes the uncoupled evolution of single-node QRC as before, but has a slightly different form than that in Eq. (2) considered in the main text, since the single-node QRC is incorporated differently in the measurement chain. More precisely,

$$\tilde{\mathcal{L}}_{\text{QRC}}\hat{\rho} = -i\left[\omega_d\hat{b}_1^\dagger\hat{b}_1 - \frac{\Lambda_1}{2}\hat{b}_1^\dagger\hat{b}_1^\dagger\hat{b}_1\hat{b}_1, \hat{\rho}\right]. \quad (\text{L3})$$

Next, \mathcal{L}_{PA} describes the evolution of the PA, defined as a non-degenerate (phase-preserving) parametric amplifier with interaction strength G_{PA} ,

$$\begin{aligned} \mathcal{L}_{\text{PA}}\hat{\rho} = & -i\left[\omega_{d1}\hat{d}_1^\dagger\hat{d}_1 + \omega_{d2}\hat{d}_2^\dagger\hat{d}_2 + (G_{\text{PA}}\hat{d}_1\hat{d}_2 + h.c.), \hat{\rho}\right] \\ & + \gamma_{d1}\mathcal{D}[\hat{d}_1]\hat{\rho} + \gamma_{d2}\mathcal{D}[\hat{d}_2]\hat{\rho}, \end{aligned} \quad (\text{L4})$$

with frequencies ω_{d1}, ω_{d2} and damping rates γ_{d1}, γ_{d2} for the two PA modes. Finally, the coupling Liouvillian $\tilde{\mathcal{L}}_c$ between the single nonlinear node of the QRC and mode \hat{d}_1 of the PA takes the form

$$\tilde{\mathcal{L}}_c\hat{\rho} = -i\left[\frac{\tilde{g}_c}{2}\hat{b}_1^\dagger\hat{d}_1 + h.c., \hat{\rho}\right] + \tilde{\Gamma}_c\mathcal{D}[\hat{b}_1 + \hat{d}_1]\hat{\rho}, \quad (\text{L5})$$

which is exactly the same form as the coupling between the two-mode amplifier system and the QRC, and represents a circulator.

Finally, we now consider heterodyne measurements of a single node of the PA, in this case \hat{d}_1 , so that the stochastic measurement superoperator $\tilde{\mathcal{S}}_{\text{meas}}(dW)$ takes the form

$$\begin{aligned} \tilde{\mathcal{S}}_{\text{meas}}(dW)\hat{\rho}^c = & \sqrt{\frac{\gamma_{d1}}{2}}\left(\hat{d}_1\hat{\rho}^c + \hat{\rho}^c\hat{d}_1^\dagger - \langle\hat{d}_1 + \hat{d}_1^\dagger\rangle^c\right)dW_d^X(t) + \\ & \sqrt{\frac{\gamma_{d1}}{2}}\left(-i\hat{d}_1\hat{\rho}^c + i\hat{\rho}^c\hat{d}_1^\dagger - \langle-i\hat{d}_1 + i\hat{d}_1^\dagger\rangle^c\right)dW_d^P(t) \end{aligned} \quad (\text{L6})$$

and the obtained measurement currents $J_d^{X,P}(t)$ are given by:

$$J_d^X(t) dt = \sqrt{\frac{\gamma_{d1}}{2}}\langle\hat{d}_1 + \hat{d}_1^\dagger\rangle^c dt + dW_d^X(t), \quad (\text{L7})$$

$$J_d^P(t) dt = \sqrt{\frac{\gamma_{d1}}{2}}\langle-i\hat{d}_1 + i\hat{d}_1^\dagger\rangle^c dt + dW_d^P(t). \quad (\text{L8})$$

Having defined the SME for the measurement chain depicted in Fig. 16(a), we now discuss the choice of parameters to operate the measurement chain in the regime of interest. We first write down the unconditional equations of motion for first-order moments of the single nonlinear node \hat{b}_1 and the

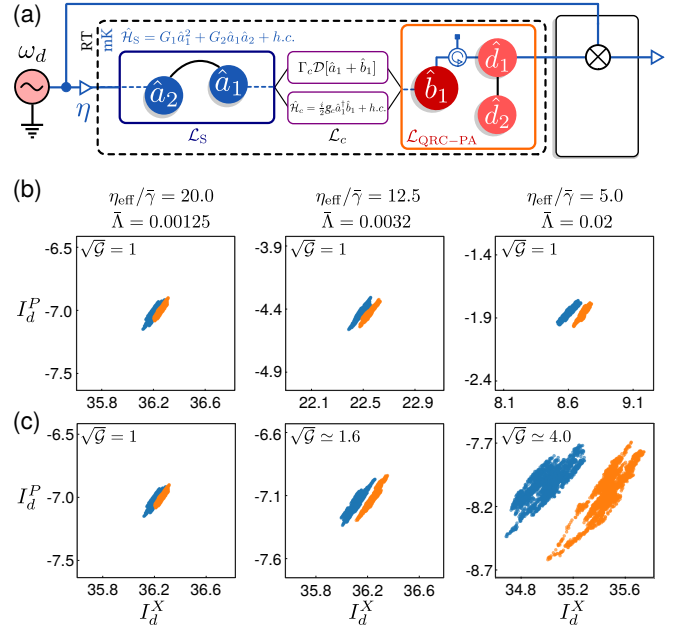


Figure 16. (a) Schematic using a non-degenerate post-amplifier (PA) to amplify output from the single nonlinear QRC node processing inputs from a measured quantum system. (b) Measured phase space outputs for unit gain $\sqrt{\mathcal{G}} = 1$ for different instances of the amplifier state classification task, carried out using three QRCs with $\mathcal{N}_{\text{eff}} \simeq 0.385$. Equal areas of measured phase space are shown; note, however, the different regions, as the amplitudes of measured heterodyne records are different due to the different η_{eff} strengths. (c) Measured phase space outputs with PA gain $\sqrt{\mathcal{G}}$ (as indicated) chosen to yield approximately equal output amplitudes. Again, equal areas of phase space are shown, now also in approximately the same region of phase space, as the PA appropriately amplifies the QRC output for the different η_{eff} strengths.

amplifier mode \hat{d}_1 in the frame rotating with the drive:

$$\begin{aligned} \frac{d}{dt}\langle\hat{b}_1\rangle = & \text{tr}\{\mathcal{L}_{\text{QRC}}\hat{\rho}\hat{b}_1\} + \text{tr}\{\mathcal{L}_c\hat{\rho}\hat{b}_1\} \\ = & \text{tr}\left\{-i\left[\omega_d\hat{b}_1^\dagger\hat{b}_1 - \frac{\Lambda_1}{2}\hat{b}_1^\dagger\hat{b}_1^\dagger\hat{b}_1\hat{b}_1, \hat{\rho}\right]\hat{b}_1\right\} \\ & - \frac{\Gamma_c + \tilde{\Gamma}_c}{2}\langle\hat{b}_1\rangle + \frac{\tilde{g}_c - \tilde{\Gamma}_c}{2}\langle\hat{d}_1\rangle \end{aligned} \quad (\text{L9})$$

$$\begin{aligned} \frac{d}{dt}\langle\hat{d}_1\rangle = & \text{tr}\{\mathcal{L}_{\text{PA}}\hat{\rho}\hat{d}_1\} + \text{tr}\{\mathcal{L}_c\hat{\rho}\hat{d}_1\} \\ = & \left(i\Delta_{d1} - \frac{\gamma_{d1} + \tilde{\Gamma}_c}{2}\right)\langle\hat{d}_1\rangle - iG_{\text{PA}}\langle\hat{d}_2^\dagger\rangle \\ & + \frac{\tilde{g}_c + \tilde{\Gamma}_c}{2}\langle\hat{b}_1\rangle \end{aligned} \quad (\text{L10})$$

$$\begin{aligned} \frac{d}{dt}\langle\hat{d}_2\rangle = & \text{tr}\{\mathcal{L}_{\text{PA}}\hat{\rho}\hat{d}_2\} + \text{tr}\{\mathcal{L}_c\hat{\rho}\hat{d}_2\} \\ = & \left(i\Delta_{d2} - \frac{\gamma_{d2}}{2}\right)\langle\hat{d}_2\rangle - iG_{12}\langle\hat{d}_1^\dagger\rangle \end{aligned} \quad (\text{L11})$$

where we have introduced the detunings $\Delta_{d1,d2} = \omega_d - \omega_{d1,d2}$, which will be taken to vanish for convenience. For a directional coupling such that the QRC node field drives the

PA, but signals from the PA do not impact the unconditional dynamics of the QRC node, we once more choose

$$\tilde{g}_c = \tilde{\Gamma}_c. \quad (\text{L12})$$

As is clear from Eq. (L9), the decay rate of the QRC node is now given by $\Gamma_c + \tilde{\Gamma}_c$. Recall that the decay rate of the QRC considered in Sec. V was given by $\Gamma_c + \gamma_1 = 3\kappa/2$, with $\Gamma_c = \kappa/2$ and $\gamma_1 = \kappa$. We will require the single QRC node in this measurement chain to have the same total decay rate, namely $\Gamma_c + \tilde{\Gamma}_c = 3\kappa/2$; keeping the same value of $\Gamma_c = \kappa/2$, we therefore must have $\tilde{\Gamma}_c = \kappa$. For convenience, we also enforce both PA modes to have the same total decay rate, so that $\gamma_{d1} + \tilde{\Gamma}_c = \gamma_{d2} \equiv \gamma_d$. We do so by requiring $\gamma_{d1} = \kappa/2, \gamma_{d2} = 3\kappa/2$.

With these parameter choices, we can proceed to analyze the dynamics of the PA nodes when driven by signals from the single nonlinear QRC node. In particular, we can calculate the gain by which the post-amplifier enhances the QRC signal. This can be extracted simply from the equations of motion of the first-order moments of the post-amplifier, written for convenience in the quadrature basis,

$$\frac{d}{dt} \begin{pmatrix} \langle \hat{X}_{d1} \rangle \\ \langle \hat{P}_{d1} \rangle \\ \langle \hat{P}_{d2} \rangle \\ \langle \hat{X}_{d2} \rangle \end{pmatrix} = \tilde{\mathbf{M}} \begin{pmatrix} \langle \hat{X}_{d1} \rangle \\ \langle \hat{P}_{d1} \rangle \\ \langle \hat{X}_{d2} \rangle \\ \langle \hat{P}_{d2} \rangle \end{pmatrix} + \Gamma'_c \begin{pmatrix} \langle \hat{X}_1 \rangle \\ \langle \hat{P}_1 \rangle \\ 0 \\ 0 \end{pmatrix} \quad (\text{L13})$$

where the dynamical matrix $\tilde{\mathbf{M}}$ takes the form:

$$\tilde{\mathbf{M}} = \begin{pmatrix} -\frac{\gamma_d}{2} & 0 & 0 & -G_{\text{PA}} \\ 0 & -\frac{\gamma_d}{2} & -G_{\text{PA}} & 0 \\ 0 & -G_{\text{PA}} & -\frac{\gamma_d}{2} & 0 \\ -G_{\text{PA}} & 0 & 0 & -\frac{\gamma_d}{2} \end{pmatrix} \quad (\text{L14})$$

Solving for the steady-state first-order moments of the PA mode \hat{d}_1 , we immediately find

$$\langle \hat{X}_{d1} \rangle = -\frac{2\gamma_d}{4G_{\text{PA}}^2 - \gamma_d^2} \cdot \tilde{\Gamma}_c \langle \hat{X}_1 \rangle. \quad (\text{L15})$$

To define the transmission gain \mathcal{G} of the PA, we can now simply write down the ratio of its output field amplitude, defined via the ensemble-averaged measurement record $\langle J_d^X \rangle$, relative to the input to be amplified, namely the single QRC node field:

$$\frac{\langle J_d^X \rangle}{\sqrt{\tilde{\Gamma}_c} \langle \hat{X}_1 \rangle} = \frac{\sqrt{\gamma_{d1}} \langle \hat{X}_{d1} \rangle}{\sqrt{\tilde{\Gamma}_c} \langle \hat{X}_1 \rangle} = \left[-\frac{2\gamma_d \sqrt{\tilde{\Gamma}_c} \gamma_{d1}}{4G_{\text{PA}}^2 - \gamma_d^2} \right] \equiv \sqrt{\mathcal{G}} \quad (\text{L16})$$

For $G_{\text{PA}} < \gamma_d/2$ (the instability threshold for the PA), the factor in square brackets is positive-definite. The same gain applies to the heterodyne measurement of the conjugate quadrature $\langle \hat{J}_d^P \rangle$ due to the non-degenerate (i.e. phase-preserving) nature of the PA chosen here.

The required PA gain determines the interaction strength G_{PA} , thus specifying the remaining operating parameters

of the measurement chain for a given task. We will apply this measurement chain to the task of classifying amplifier states discussed in Sec. V of the main text, in particular when using a single-node QRC to classify states with different values of η_{eff} , as analyzed in Sec. VD. More precisely, we consider three instances of the task as before, defined by effective drive strengths $\eta_{\text{eff}}/\bar{\gamma} \in \{5.0, 12.5, 20.0\}$. For each instance, the nonlinearity of the single-node QRC in the measurement chain is respectively chosen as $\bar{\Lambda}/\bar{\gamma} \in \{0.00125, 0.0032, 0.02\}$, to render the effective dimensionless nonlinearity parameter \mathcal{N}_{eff} equal for QRCs carrying out the corresponding task; for the choice of parameters here, $\mathcal{N}_{\text{eff}} \simeq 0.385$.

We first consider the case where the PA is operated to provide unit gain. This requires the factor in Eq. (L16) to have a magnitude of one, which imposes the PA interaction strength:

$$G_{\text{PA}}^{\mathcal{G}=1} = \sqrt{\frac{1}{4}\gamma_d^2 - \frac{1}{2}\gamma_d\sqrt{\tilde{\Gamma}_c}\gamma_{d1}} = \frac{\kappa}{4}\sqrt{9 - 6\sqrt{2}} \simeq 0.1794\kappa. \quad (\text{L17})$$

The measured QRC-PA output is constructed as before using measured quadratures obtained from the PA, defined analogously to Sec. V,

$$\mathbf{x}(t) = \begin{pmatrix} I_d^X(t) \\ I_d^P(t) \end{pmatrix}, \quad I_d^{X,P}(t) = \left\langle \frac{1}{t} \int_{t_0}^t d\tau J_d^{X,P}(\tau) \right\rangle_{N_S}, \quad (\text{L18})$$

and are plotted in Fig. 16(b) for the three instances of the classification task from left to right using QRCs with $\mathcal{N}_{\text{eff}} \text{ simeq } 0.385$, and for $N_S = 200$. Due to the directionality of the coupling between the QRC and the PA, for unit gain operation the output of the measurement chain including the PA is quite similar to that without the PA, as considered in Sec. V. Hence the salient features of classification performance for different η_{eff} are the same as those observed in Fig. 11. In particular, output from the measured chain with the QRC possessing the strongest nonlinearity (rightmost panel) has the largest separation of measured distributions. However, as noted earlier, this output also has the lowest amplitudes of measured quadratures, as evident from the phase space axes, due to the weakest input drive strength.

The observed difference can be overcome by choosing the PA interaction strength G_{PA} such that for all three instances of the classification task, the amplitude of measured outputs is approximately equal. For concreteness, we do so by requiring the measured output amplitudes for the weaker effective drives $\eta_{\text{eff}}/\bar{\gamma} \in \{5.0, 12.5\}$, to match the measured output amplitudes for the strongest effective drive, $\eta_{\text{eff}}/\bar{\gamma} = 20.0$. This requires amplitude gains $\sqrt{\mathcal{G}}$ determined (approximately) by the ratio of the corresponding effective drives, $\sqrt{\mathcal{G}} \in \{20.0/5.0 = 4.0, 20.0/12.5 = 1.6, 1.0\}$ respectively for the three instances considered.

Applying this PA gain, the measured quadrature distributions are shown in Fig. 16(c). The measured output distributions are now in very similar regions of phase space, indicating similar magnitudes of measured outputs. The lack of an exact equivalence is due to the change in PA linewidth for

increasing gain due to anti-damping, which modifies the PA response time. As a result, for larger gain, the original value of waiting time t_0 introduced in Eq. (36) and carried through here is no longer sufficient for the measurement chain output to settle to steady-state. Consequently, the observed gain is a slight underestimate of the steady-state gain \sqrt{G} , for large gains.

Each panel in Figs. 16(b), (c) show the same total area in phase space, albeit centered around different points. We can thus compare the variance of distributions directly. When

comparing each measured distribution to its counterpart at unit PA gain (the panel directly above), we clearly see that phase-preserving amplification for $G > 1$ adds amplified quantum fluctuations to both measured quadratures, leading to an increase in the variance of measured distributions. This is particularly clear for the measurement chain with weakest input drive, which requires the largest gain amplification. However, the displacement of distributions that enables classification is also enhanced. As a result, the measurement chain containing the QRC with strongest nonlinearity would still require fewer experimental resources for classification.

-
- [1] H. Jaeger, Harnessing nonlinearity: Predicting chaotic systems and saving energy in wireless communication, *Science* **304**, 78–80 (2004).
- [2] M. Lukoševičius and H. Jaeger, Reservoir computing approaches to recurrent neural network training, *Computer Science Review* **3**, 127 (2009).
- [3] G. Tanaka, T. Yamane, J. B. Héroux, R. Nakane, N. Kanazawa, S. Takeda, H. Numata, D. Nakano, and A. Hirose, Recent advances in physical reservoir computing: A review, *Neural Networks* **115**, 100–123 (2019).
- [4] D. J. Gauthier, E. Bollt, A. Griffith, and W. A. S. Barbosa, Next generation reservoir computing, *Nature Communications* **12**, 10.1038/s41467-021-25801-2 (2021).
- [5] J. Pathak, B. Hunt, M. Girvan, Z. Lu, and E. Ott, Model-Free Prediction of Large Spatiotemporally Chaotic Systems from Data: A Reservoir Computing Approach, *Physical Review Letters* **120**, 024102 (2018).
- [6] A. Griffith, A. Pomerance, and D. J. Gauthier, Forecasting chaotic systems with very low connectivity reservoir computers, *Chaos: An Interdisciplinary Journal of Nonlinear Science* **29**, 123108 (2019).
- [7] D. J. Gauthier and I. Fischer, Predicting hidden structure in dynamical systems, *Nature Machine Intelligence* **3**, 281 (2021).
- [8] K. Nakai and Y. Saiki, Machine-learning inference of fluid variables from data using reservoir computing, *Physical Review E* **98**, 023111 (2018).
- [9] A. Röhm, D. J. Gauthier, and I. Fischer, Model-free inference of unseen attractors: Reconstructing phase space features from a single noisy trajectory using reservoir computing, [arXiv:2108.04074 \[nlin\]](https://arxiv.org/abs/2108.04074) (2021).
- [10] D. Canaday, A. Pomerance, and D. J. Gauthier, Model-free control of dynamical systems with deep reservoir computing, *Journal of Physics: Complexity* **2**, 035025 (2021).
- [11] S. L. Brunton, J. L. Proctor, and J. N. Kutz, Discovering governing equations from data by sparse identification of nonlinear dynamical systems, *Proceedings of the National Academy of Sciences* **113**, 3932 (2016).
- [12] J. Pathak, Z. Lu, B. R. Hunt, M. Girvan, and E. Ott, Using machine learning to replicate chaotic attractors and calculate Lyapunov exponents from data, *Chaos: An Interdisciplinary Journal of Nonlinear Science* **27**, 121102 (2017).
- [13] L. Larger, A. Baylón-Fuentes, R. Martinenghi, V. S. Udaltsov, Y. K. Chembo, and M. Jacquot, High-Speed Photonic Reservoir Computing Using a Time-Delay-Based Architecture: Million Words per Second Classification, *Physical Review X* **7**, 011015 (2017).
- [14] M. Rafayelyan, J. Dong, Y. Tan, F. Krzakala, and S. Gigan, Large-Scale Optical Reservoir Computing for Spatiotemporal Chaotic Systems Prediction, *Physical Review X* **10**, 041037 (2020), publisher: American Physical Society.
- [15] W. A. S. Barbosa, A. Griffith, G. E. Rowlands, L. C. G. Govia, G. J. Ribeill, M.-H. Nguyen, T. A. Ohki, and D. J. Gauthier, Symmetry-aware reservoir computing, *Physical Review E* **104**, 045307 (2021).
- [16] L. Grigoryeva and J.-P. Ortega, Echo state networks are universal, *Neural Networks* **108**, 495 (2018).
- [17] L. Gonon and J.-P. Ortega, Reservoir Computing Universality With Stochastic Inputs, *IEEE Transactions on Neural Networks and Learning Systems* **31**, 100 (2020).
- [18] 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, *Communications Physics* **4**, 1 (2021), bandiera_abtest: a Cc_license_type: cc_by Cg_type: Nature Research Journals Number: 1 Primary_atype: Research Publisher: Nature Publishing Group Subject_term: Quantum mechanics;Quantum optics Subject_term_id: quantum-mechanics;quantum-optics.
- [19] L. Larger, M. C. Soriano, D. Brunner, L. Appeltant, J. M. Gutiérrez, L. Pesquera, C. R. Mirasso, and I. Fischer, Photonic information processing beyond Turing: an optoelectronic implementation of reservoir computing, *Optics Express* **20**, 3241 (2012).
- [20] D. Brunner, M. C. Soriano, C. R. Mirasso, and I. Fischer, Parallel photonic information processing at gigabyte per second data rates using transient states, *Nature Communications* **4**, 1364 (2013), bandiera_abtest: a Cg_type: Nature Research Journals Number: 1 Primary_atype: Research Publisher: Nature Publishing Group Subject_term: Applied optics;Nonlinear optics Subject_term_id: applied-optics;nonlinear-optics.
- [21] F. Duport, B. Schneider, A. Smerieri, M. Haelterman, and S. Massar, All-optical reservoir computing, *Optics Express* **20**, 22783 (2012).
- [22] C. Mesaritakis, A. Bogris, A. Kapsalis, and D. Syvridis, High-speed all-optical pattern recognition of dispersive Fourier images through a photonic reservoir computing subsystem, *Optics Letters* **40**, 3416 (2015).
- [23] F. D.-I. Coarer, M. Sciamanna, A. Katumba, M. Freiburger, J. Dambre, P. Bienstman, and D. Rontani, All-Optical Reservoir Computing on a Photonic Chip Using Silicon-Based Ring Resonators, *IEEE Journal of Selected Topics in Quantum Electronics* **24**, 1 (2018).
- [24] J. Dong, M. Rafayelyan, F. Krzakala, and S. Gigan, Optical reservoir computing using multiple light scattering for chaotic systems prediction, *IEEE Journal of Selected Topics in Quantum Electronics* **26**, 1–12 (2020).
- [25] M. C. Soriano, S. Ortín, L. Keuninckx, L. Appeltant, J. Danck-

- aert, L. Pesquera, and G. van der Sande, Delay-Based Reservoir Computing: Noise Effects in a Combined Analog and Digital Implementation, *IEEE Transactions on Neural Networks and Learning Systems* **26**, 388 (2015).
- [26] N. D. Haynes, M. C. Soriano, D. P. Rosin, I. Fischer, and D. J. Gauthier, Reservoir computing with a single time-delay autonomous Boolean node, *Physical Review E* **91**, 020801 (2015).
- [27] D. Canaday, A. Griffith, and D. J. Gauthier, Rapid time series prediction with a hardware-based reservoir computer, *Chaos: An Interdisciplinary Journal of Nonlinear Science* **28**, 123119 (2018).
- [28] S. Apostel, N. D. Haynes, E. Schöll, O. D’Huys, and D. J. Gauthier, Reservoir Computing Using Autonomous Boolean Networks Realized on Field-Programmable Gate Arrays, in *Reservoir Computing: Theory, Physical Implementations, and Applications*, Natural Computing Series, edited by K. Nakajima and I. Fischer (Springer, Singapore, 2021) pp. 239–271.
- [29] G. E. Rowlands, M.-H. Nguyen, G. J. Ribeill, A. P. Wagner, L. C. G. Govia, W. A. S. Barbosa, D. J. Gauthier, and T. A. Ohki, Reservoir Computing with Superconducting Electronics, [arXiv:2103.02522 \[cond-mat, physics:physics\]](https://arxiv.org/abs/2103.02522) (2021).
- [30] G. Angelatos, S. Khan, and H. E. Türeci, Reservoir Computing Approach to Quantum State Measurement, [arXiv:2011.09652 \[cond-mat, physics:quant-ph\]](https://arxiv.org/abs/2011.09652) (2021).
- [31] K. Nakajima, Physical reservoir computing—an introductory perspective, *Japanese Journal of Applied Physics* **59**, 060501 (2020).
- [32] L. G. Wright, T. Onodera, M. M. Stein, T. Wang, D. T. Schachter, Z. Hu, and P. L. McMahon, Deep physical neural networks enabled by a backpropagation algorithm for arbitrary physical systems, [arXiv:2104.13386 \[cond-mat, physics:physics\]](https://arxiv.org/abs/2104.13386) (2021).
- [33] K. Nakajima and I. Fischer, *Reservoir Computing: Theory, Physical Implementations, and Applications* (Springer Nature, 2021).
- [34] K. Fujii and K. Nakajima, Harnessing disordered-ensemble quantum dynamics for machine learning, *Physical Review Applied* **8**, 024030 (2017).
- [35] S. Ghosh, A. Opala, M. Matuszewski, T. Paterek, and T. C. H. Liew, Quantum reservoir processing, *npj Quantum Information* **5**, 1 (2019).
- [36] L. G. Wright and P. L. McMahon, The capacity of quantum neural networks, [arXiv:1908.01364 \[physics, physics:quant-ph\]](https://arxiv.org/abs/1908.01364) (2019).
- [37] J. Nokkala, R. Martínez-Peña, G. L. Giorgi, V. Parigi, M. C. Soriano, and R. Zambrini, Gaussian states provide universal and versatile quantum reservoir computing, [arXiv:2006.04821 \[quant-ph\]](https://arxiv.org/abs/2006.04821) (2020).
- [38] 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, *Physical Review Research* **3**, 013077 (2021).
- [39] R. Martínez-Peña, G. L. Giorgi, J. Nokkala, M. C. Soriano, and R. Zambrini, Dynamical Phase Transitions in Quantum Reservoir Computing, *Physical Review Letters* **127**, 100502 (2021).
- [40] W. D. Kalfus, G. J. Ribeill, G. E. Rowlands, H. K. Krovi, T. A. Ohki, and L. C. G. Govia, Neuromorphic computing with a single qudit, [arXiv:2101.11729 \[quant-ph\]](https://arxiv.org/abs/2101.11729) (2021).
- [41] A. Wallraff, D. I. Schuster, A. Blais, L. Frunzio, R.-S. Huang, J. Majer, S. Kumar, S. M. Girvin, and R. J. Schoelkopf, Strong coupling of a single photon to a superconducting qubit using circuit quantum electrodynamics, *Nature* **431**, 162 (2004), bandiera_abtest: a Cg_type: Nature Research Journals Number: 7005 Primary_atype: Research Publisher: Nature Publishing Group.
- [42] A. Blais, A. L. Grimsmo, S. Girvin, and A. Wallraff, Circuit quantum electrodynamics, *Reviews of Modern Physics* **93**, 025005 (2021), publisher: American Physical Society.
- [43] H. J. Carmichael, Quantum trajectory theory for cascaded open systems, *Physical Review Letters* **70**, 2273 (1993).
- [44] H. M. Wiseman and G. J. Milburn, Quantum theory of optical feedback via homodyne detection, *Physical Review Letters* **70**, 548 (1993).
- [45] H. M. Wiseman and G. J. Milburn, Interpretation of quantum jump and diffusion processes illustrated on the Bloch sphere, *Physical Review A* **47**, 1652 (1993).
- [46] J. K. Breslin and G. J. Milburn, Conditional variance reduction by measurements on correlated field modes, *Physical Review A* **55**, 1430 (1997), publisher: American Physical Society.
- [47] A. C. Doherty and K. Jacobs, Feedback control of quantum systems using continuous state estimation, *Physical Review A* **60**, 2700 (1999), publisher: American Physical Society.
- [48] M. Boissonneault, J. M. Gambetta, and A. Blais, Dispersive regime of circuit QED: Photon-dependent qubit dephasing and relaxation rates, *Physical Review A* **79**, 013819 (2009), publisher: American Physical Society.
- [49] C. W. Sandbo Chang, M. Simoen, J. Aumentado, C. Sabín, P. Forn-Díaz, A. M. Vadiraj, F. Quijandría, G. Johansson, I. Fuentes, and C. M. Wilson, Generating Multimode Entangled Microwaves with a Superconducting Parametric Cavity, *Physical Review Applied* **10**, 044019 (2018), publisher: American Physical Society.
- [50] C. H. Bennett, D. P. DiVincenzo, J. A. Smolin, and W. K. Wootters, Mixed-state entanglement and quantum error correction, *Physical Review A* **54**, 3824 (1996), publisher: American Physical Society.
- [51] M. C. Arnesen, S. Bose, and V. Vedral, Natural Thermal and Magnetic Entanglement in the 1D Heisenberg Model, *Physical Review Letters* **87**, 017901 (2001), publisher: American Physical Society.
- [52] J. Clarke and A. I. Braginski, *The SQUID Handbook* (Wiley-VCH, Weinheim, 2003).
- [53] J. Koch, T. M. Yu, J. Gambetta, A. A. Houck, D. I. 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, *Physical Review A* **76**, 042319 (2007), publisher: American Physical Society.
- [54] H. J. Carmichael, *Statistical Methods in Quantum Optics I - Master Equations and Fokker-Planck Equations | Howard J. Carmichael | Springer* (2002).
- [55] S. Boutin, D. M. Toyli, A. V. Venkatramani, A. W. Eddins, I. Siddiqi, and A. Blais, Effect of Higher-Order Nonlinearities on Amplification and Squeezing in Josephson Parametric Amplifiers, *Physical Review Applied* **8**, 054030 (2017).
- [56] Y.-S. Ra, A. Dufour, M. Walschaers, C. Jacquard, T. Michel, C. Fabre, and N. Treps, Non-Gaussian quantum states of a multimode light field, *Nature Physics* **16**, 144 (2020).
- [57] P. D. Drummond and D. F. Walls, Quantum theory of optical bistability. I. Nonlinear polarisability model, *Journal of Physics A: Mathematical and General* **13**, 725 (1980).
- [58] M. I. Dykman and M. A. Krivoglaz, Theory of fluctuational transitions between stable states of a nonlinear oscillator, *Soviet Journal of Experimental and Theoretical Physics* **77**, 30 (1979).
- [59] M. Dykman, *Fluctuating Nonlinear Oscillators: From Nanomechanics to Quantum Superconducting Circuits* (OUP Oxford, 2012).
- [60] O. Černotík, D. V. Vasilyev, and K. Hammerer, Adiabatic elimination of Gaussian subsystems from quantum dynamics under continuous measurement, *Physical Review A* **92**, 012124

- (2015), publisher: American Physical Society.
- [61] J. Zhang and K. Mølmer, Prediction and retrodiction with continuously monitored Gaussian states, *Physical Review A* **96**, 062131 (2017).
- [62] J. R. Johansson, P. D. Nation, and F. Nori, QuTiP 2: A Python framework for the dynamics of open quantum systems, *Computer Physics Communications* **184**, 1234 (2013).
- [63] S. Filipp, P. Maurer, P. J. Leek, M. Baur, R. Bianchetti, J. M. Fink, M. Göppl, L. Steffen, J. M. Gambetta, A. Blais, and A. Wallraff, Two-Qubit State Tomography Using a Joint Dispersive Readout, *Physical Review Letters* **102**, 200402 (2009), publisher: American Physical Society.
- [64] A. Metelmann and A. A. Clerk, Nonreciprocal photon transmission and amplification via reservoir engineering, *Phys. Rev. X* **5**, 021025 (2015).
- [65] K. Sliwa, M. Hatridge, A. Narla, S. Shankar, L. Frunzio, R. Schoelkopf, and M. Devoret, Reconfigurable Josephson Circulator/Directional Amplifier, *Physical Review X* **5**, 041020 (2015), publisher: American Physical Society.
- [66] T.-C. Chien, O. Lanes, C. Liu, X. Cao, P. Lu, S. Motz, G. Liu, D. Pekker, and M. Hatridge, Multiparametric amplification and qubit measurement with a Kerr-free Josephson ring modulator, *Physical Review A* **101**, 042336 (2020), publisher: American Physical Society.
- [67] 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).
- [68] K. W. Murch, S. J. Weber, C. Macklin, and I. Siddiqi, Observing single quantum trajectories of a superconducting quantum bit, *Nature* **502**, 211 (2013).
- [69] T. Walter, P. Kurpiers, S. Gasparinetti, P. Magnard, A. Potočnik, Y. Salathé, M. Pechal, M. Mondal, M. Oppliger, C. Eichler, and A. Wallraff, Rapid high-fidelity single-shot dispersive readout of superconducting qubits, *Phys. Rev. Applied* **7**, 054020 (2017).
- [70] L. S. Bishop, Circuit Quantum Electrodynamics, [arXiv:1007.3520](https://arxiv.org/abs/1007.3520) [cond-mat, physics:quant-ph] (2010), [arXiv: 1007.3520](https://arxiv.org/abs/1007.3520).
- [71] M. Hatridge, R. Vijay, D. H. Slichter, J. Clarke, and I. Siddiqi, Dispersive magnetometry with a quantum limited SQUID parametric amplifier, *Physical Review B* **83**, 134501 (2011), publisher: American Physical Society.
- [72] C. Macklin, K. O'Brien, D. Hover, M. E. Schwartz, V. Bolkhovskiy, X. Zhang, W. D. Oliver, and I. Siddiqi, A near-quantum-limited Josephson traveling-wave parametric amplifier, *Science* **350**, 307 (2015), publisher: American Association for the Advancement of Science.
- [73] A. Roy and M. Devoret, Introduction to parametric amplification of quantum signals with Josephson circuits, *Comptes Rendus Physique Quantum microwaves / Micro-ondes quantiques*, **17**, 740 (2016).
- [74] F. Arute, K. Arya, R. Babbush, D. Bacon, J. C. Bardin, R. Barends, R. Biswas, S. Boixo, F. G. S. L. Brandao, D. A. Buell, B. Burkett, Y. Chen, Z. Chen, B. Chiaro, R. Collins, W. Courtney, A. Dunsworth, E. Farhi, B. Foxen, A. Fowler, C. Gidney, M. Giustina, R. Graff, K. Guerin, S. Habegger, M. P. Harrigan, M. J. Hartmann, A. Ho, M. Hoffmann, T. Huang, T. S. Humble, S. V. Isakov, E. Jeffrey, Z. Jiang, D. Kafri, K. Kechedzhi, J. Kelly, P. V. Klimov, S. Knysh, A. Korotkov, F. Kostritsa, D. Landhuis, M. Lindmark, E. Lucero, D. Lyakh, S. Mandrà, J. R. McClean, M. McEwen, A. Megrant, X. Mi, K. Michielsen, M. Mohseni, J. Mutus, O. Naaman, M. Neeley, C. Neill, M. Y. Niu, E. Ostby, A. Petukhov, J. C. Platt, C. Quintana, E. G. Rieffel, P. Roushan, N. C. Rubin, D. Sank, K. J. Satzinger, V. Smelyanskiy, K. J. Sung, M. D. Trevithick, A. Vainsencher, B. Villalonga, T. White, Z. J. Yao, P. Yeh, A. Zalcman, H. Neven, and J. M. Martinis, Quantum supremacy using a programmable superconducting processor, *Nature* **574**, 505 (2019).
- [75] Y. Wu, W.-S. Bao, S. Cao, F. Chen, M.-C. Chen, X. Chen, T.-H. Chung, H. Deng, Y. Du, D. Fan, M. Gong, C. Guo, C. Guo, S. Guo, L. Han, L. Hong, H.-L. Huang, Y.-H. Huo, L. Li, N. Li, S. Li, Y. Li, F. Liang, C. Lin, J. Lin, H. Qian, D. Qiao, H. Rong, H. Su, L. Sun, L. Wang, S. Wang, D. Wu, Y. Xu, K. Yan, W. Yang, Y. Yang, Y. Ye, J. Yin, C. Ying, J. Yu, C. Zha, C. Zhang, H. Zhang, K. Zhang, Y. Zhang, H. Zhao, Y. Zhao, L. Zhou, Q. Zhu, C.-Y. Lu, C.-Z. Peng, X. Zhu, and J.-W. Pan, Strong Quantum Computational Advantage Using a Superconducting Quantum Processor, *Physical Review Letters* **127**, 180501 (2021), publisher: American Physical Society.
- [76] C. Berke, E. Varvelis, S. Trebst, A. Altland, and D. P. DiVincenzo, Transmon platform for quantum computing challenged by chaotic fluctuations, [arXiv:2012.05923](https://arxiv.org/abs/2012.05923) [cond-mat, physics:quant-ph] (2020), [arXiv: 2012.05923](https://arxiv.org/abs/2012.05923).
- [77] E. Bollt, On explaining the surprising success of reservoir computing forecaster of chaos? the universal machine learning dynamical system with contrast to var and dmd, *Chaos: An Interdisciplinary Journal of Nonlinear Science* **31**, 013108 (2021).
- [78] R. J. Glauber, Coherent and Incoherent States of the Radiation Field, *Physical Review* **131**, 2766 (1963).
- [79] E. C. G. Sudarshan, Equivalence of Semiclassical and Quantum Mechanical Descriptions of Statistical Light Beams, *Physical Review Letters* **10**, 277 (1963).
- [80] C. Gardiner, *Stochastic Methods: A Handbook for the Natural and Social Sciences*, 4th ed., Springer Series in Synergetics (Springer-Verlag, Berlin Heidelberg, 2009).