

Quantum reservoir computing with a single nonlinear oscillator

L. C. G. Govia¹*, G. J. Ribeill, G. E. Rowlands¹, H. K. Krovi, and T. A. Ohki¹
Raytheon BBN Technologies, 10 Moulton St., Cambridge, Massachusetts 02138, USA



(Received 7 May 2020; accepted 7 January 2021; published 25 January 2021)

Realizing the promise of quantum information processing remains a daunting task given the omnipresence of noise and error. Adapting noise-resilient classical computing modalities to quantum mechanics may be a viable path towards near-term applications in the noisy intermediate-scale quantum era. Here, we propose continuous variable quantum reservoir computing in a single nonlinear oscillator. Through numerical simulation of our model we demonstrate quantum-classical performance improvement and identify its likely source: the nonlinearity of quantum measurement. Beyond quantum reservoir computing, this result may impact the interpretation of results across quantum machine learning. We study how the performance of our quantum reservoir depends on Hilbert space dimension, how it is impacted by injected noise, and briefly comment on its experimental implementation. Our results show that quantum reservoir computing in a single nonlinear oscillator is an attractive modality for quantum computing on near-term hardware.

DOI: 10.1103/PhysRevResearch.3.013077

I. INTRODUCTION

Over the last several decades, quantum information science has emerged as a transformative framework for information processing, from high-performance computing to communication and cryptography [1]. Despite the tremendous potential, only recently has a computational task that is prohibitively difficult on a classical computer been demonstrated on quantum hardware [2]. Intrinsic noise and spurious error remain the roadblock to widespread quantum advantage, and the full power of quantum information processing awaits the demonstration of logical error correction and fault tolerance [3,4]. In the near term, while we remain in the noisy intermediate-scale quantum (NISQ) era [5], open questions remain as to what, if any, quantum advantage can be expected from current technologies, and how to design noise-resilient computational methodologies that best take advantage of the quantum resources available [6].

In the high-performance computing frontier, neuromorphic, or brain-inspired, computing modalities [7] show considerable promise, with the canonical example being the widespread success of neural network approaches to machine learning. A particularly attractive neuromorphic computing approach is reservoir computing [8–10], which harnesses the computational power of a disordered, sparsely-connected, nonlinear network. The network connections are untrained and remain fixed, which both drastically reduces the cost of training and removes the susceptibility to error in the assign-

ment of flexible network connections as in traditional neural networks. Recently, this approach to neuromorphic computing has been expanded to include quantum mechanical systems as the reservoir, defining so-called quantum reservoir computing [11,12]. Theoretical studies have shown application to problems in classical computing [13–17] as well as quantum computing [18,19], and first experimental demonstrations have been reported [20–22].

In this work, we present quantum reservoir computing (QRC) in a continuous-variable system, with a reservoir formed by a single nonlinear oscillator, and contrast to classical reservoir computing (CRC) with the equivalent classical reservoir. By using a continuous-variable system, we reduce the costly repetitions necessary to obtain accurate measurement of expectation values, an issue that affects the run time of discrete-variable quantum machine learning approaches [17]. We expect this to be an advantage that continuous-variable reservoirs will show over discrete-variable ones in practical implementation. We demonstrate via numerical simulation an improvement in performance of QRC compared to CRC for the same classical task of sine phase estimation. This improvement is both in average error for small training set sizes and a reduction in performance spread across reservoir parameters. We isolate the main cause of this quantum improvement, presenting compelling evidence that it is due to the intrinsic nonlinearity of quantum measurement.

This paper is organized as follows. In Sec. II we briefly review reservoir computing, and in Sec. III we introduce quantum and classical reservoir computing with a single nonlinear oscillator. We describe our chosen task in Sec. IV and present the results of our numerical simulations in Sec. V, studying the quantum improvement and its origin, as well as the effects of Hilbert space dimension and injected noise. Finally, we present our concluding remarks and a discussion of experimental implementations in Sec. VI.

*luke.c.govia@raytheon.com

II. RESERVOIR COMPUTING

The simplest implementation of a reservoir computer consists of a network of N sparsely-connected nonlinear nodes, into which data is fed in a time series manner. If the $N \times 1$ vector $\vec{u}(t)$ defines the input signal at time t , then the input to the reservoir nodes is given by $\mathbf{W}_{\text{in}}\vec{u}(t)$, where \mathbf{W}_{in} is the $N \times N$ input weight matrix. In practice, \mathbf{W}_{in} is fixed for a given task, and N can be replaced by N_{in} if data is input to only a restricted number of reservoir nodes.

Without loss of generality we are free to assume each node in the reservoir has a single degree of freedom, and the full state of the reservoir at time t is described by the vector $\vec{x}(t)$. The time evolution of the reservoir is given by the nonlinear differential equation

$$\partial_t \vec{x} = f(\mathbf{W}_{\text{res}}\vec{x}(t), \mathbf{W}_{\text{in}}\vec{u}(t), \mathbf{W}_{\text{fb}}\vec{x}(t - \tau)), \quad (1)$$

where $f(\vec{a}, \vec{b}, \vec{c})$ is a nonlinear function. Its first argument describes the internal interaction between reservoir nodes with corresponding internal weight matrix \mathbf{W}_{res} that includes decay of the state in each node. Its second describes the input signal to the reservoir, as previously discussed. The final input to f describes time-delay feedback in the reservoir interaction, with weight matrix \mathbf{W}_{fb} . In this work we will not consider any explicit time-delay feedback and rely on the internal reservoir dynamics to serve as short-term “memory.”

We label the output signal of the reservoir as $\vec{s}(t)$, which in many cases is simply the internal state $\vec{x}(t)$. In order to compute with the reservoir, we discretize the time-dependent output signals of all N reservoir nodes into P time steps and collect them into an $N \times P$ matrix \mathbf{s}_t . For our purposes we column stack this to form a $NP \times 1$ vector \vec{s}_t . The final step in a reservoir computing architecture is application of the $L \times NP$ output weight matrix to the output signal, to obtain the computed output

$$\vec{y} = \mathbf{W}_{\text{out}}\vec{s}_t, \quad (2)$$

where y is the $L \times 1$ task output vector and (ideally) is the answer to the problem the reservoir computer aims to solve. As with the input weight matrix, if the output signal is only recorded from a subset of nodes N_{out} , then the output weight matrix can be reduced to $L \times (N_{\text{out}}P)$ in size.

In reservoir computing, only the values of the output weight matrix are trained in a supervised learning fashion. We use the standard approach for training (see, e.g., Ref. [23])

$$\mathbf{W}_{\text{out}} = \mathbf{Y}\mathbf{S}_t^{\text{T}}(\mathbf{S}_t\mathbf{S}_t^{\text{T}} + \gamma\mathbb{I})^{-1}, \quad (3)$$

where \mathbf{Y} and \mathbf{S}_t are $L \times M$ and $(N_{\text{out}}P) \times M$ matrices that contain the data for M training instances of the task. \mathbf{Y} contains the correct output of the task and \mathbf{S}_t the reservoir output signals \vec{s}_t , for each training instance. Here, γ is a ridge-regression parameter used to prevent overfitting.

As the training phase for a reservoir computer requires only a single matrix inversion, it offers considerable computational savings over traditional neural networks. Note that the way we have designed our training procedure allows it to access both correlations in the output signal between nodes at a given instance of time, and across time steps, which stems from the fact that we train using the full output signal \vec{s}_t . We find that

this gives the best performance. To further reduce the computational cost of training, one can introduce a block-diagonal \mathbf{W}_{out} that does not mix reservoir output signals at different time steps.

III. A SINGLE NONLINEAR OSCILLATOR AS A RESERVOIR

An appealing aspect of reservoir computing is that, in principle, any nonlinear system can be used as the reservoir. For the model of quantum reservoir computing we consider here, the input and output from the reservoir will be a classical data stream, and the reservoir consists of a single nonlinear oscillator [23–26]. While it might appear that a single oscillator is a reservoir with only one node, this is in fact not the case. In our formalism, the number of nodes in the reservoir map to the number of independent degrees of freedom in the oscillator.

In a classical nonlinear oscillator, there are two independent degrees of freedom and thus two computational nodes; we choose these to be the position X and momentum P quadratures. In a quantum nonlinear oscillator, the computational nodes are formed by the complex amplitudes of the system state ρ expressed in some basis. Equivalently, one can think of the computational nodes as the real-valued expectation values of a complete basis of observables, such as the Fock state projectors. We will describe the computational nodes by a specification of expectation values of the form

$$E_{nm} = \text{Tr}[\rho \hat{X}^n \hat{P}^m] \quad (4)$$

for integers $n, m \geq 0$, where

$$\hat{X} = \frac{1}{\sqrt{2}}(\hat{a} + \hat{a}^{\dagger}), \quad \hat{P} = \frac{-i}{\sqrt{2}}(\hat{a} - \hat{a}^{\dagger}) \quad (5)$$

are the usual canonical quadratures, with \hat{a} the lowering operator. Each E_{nm} is an independent parameter, so given the infinite dimension of a quantum oscillator’s Hilbert space, it is tempting to assume that such a quantum reservoir is infinite in size.

However, for many realistic states of the oscillator the total number of independent E_{nm} is finite. An extreme example is Gaussian states, which are fully described by the set of expectation values with $n + m \leq 2$, but even an input power restriction will set a limit to the largest occupied Fock state, which in turn implies only a finite number of independent E_{nm} . Nevertheless, a quantum nonlinear oscillator can have more degrees of freedom than its classical analog, and the exponential growth of Hilbert space for oscillators is analogous to the arguments used to motivate the power of quantum reservoirs built from qubit networks [11].

Throughout this work, we consider a Kerr nonlinear oscillator as our reservoir. The quantum version of this system is described by the Hamiltonian (in a frame rotating at the oscillators’ frequency)

$$\hat{H}(t) = K\hat{a}^{\dagger}\hat{a}\hat{a}^{\dagger}\hat{a} + \alpha u(t)(\hat{a} + \hat{a}^{\dagger}), \quad (6)$$

where K is the Kerr nonlinearity, and α is an overall amplitude to the scalar input $u(t)$. As can be seen, we choose a linear input coupling that drives a single degree of freedom of the reservoir, namely the P quadrature. As output we measure the expectation value of the X quadrature (again only one degree

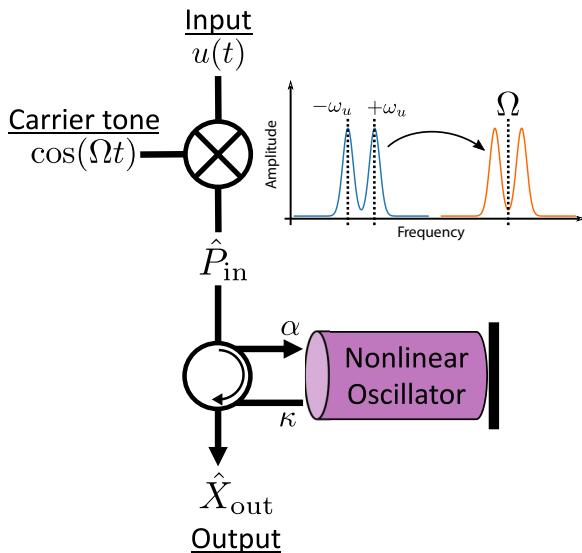


FIG. 1. Input-output scheme for nonlinear oscillator RC. A sinusoidal input signal $u(t)$ at frequency ω_u is up-converted to a double-sideband modulation of a carrier tone centered at the nonlinear oscillator resonance frequency Ω , where $\Omega \gg \omega_u$. This is input to the P quadrature of the oscillator, and the X quadrature is continuously measured by monitoring the output field of the oscillator. The circulator at the input/output port of the nonlinear oscillator separates the propagation of the input and output signals.

of freedom)

$$s(t) = E_{10} = \text{Tr}[\rho(t)\hat{X}], \quad (7)$$

as a function of time. At the beginning of each new input to the nonlinear oscillator, we reinitialize the system in a deterministic state, in this case vacuum.

The input-output diagram for our setup is shown in Fig. 1. Under driving by an input field in the P quadrature at a carrier frequency resonant with the zero-photon frequency of the nonlinear oscillator, labeled \hat{P}_{in} in Fig. 1, the dynamical evolution of the internal degrees of freedom of the nonlinear oscillator is generated by the Hamiltonian of Eq. (6). Using input-output theory [27], the output field in the X quadrature, labeled \hat{X}_{out} in Fig. 1, is given by

$$\hat{X}_{\text{out}}(t) = \hat{X}_{\text{in}}(t) + \sqrt{\kappa}\hat{X}(t). \quad (8)$$

Since the input is entirely in the P quadrature, we have that $\langle \hat{X}_{\text{in}}(t) \rangle = 0$, and so the expectation value of the output field is given by

$$\langle \hat{X}_{\text{out}}(t) \rangle = \sqrt{\kappa}\langle \hat{X}(t) \rangle = s(t), \quad (9)$$

which is exactly the signal we have chosen as our reservoir output. Thus, our chosen signal can be continuously measured *in practice* by monitoring the output field of the nonlinear oscillator.

In simulation, we can directly access the internal degrees of freedom and evaluate $s(t)$. However, we need to account for the impact of the experimentally relevant open-system measurement setup on the system dynamics. To do so, we model the evolution of the quantum reservoir using a Lindblad

master equation [28]

$$\dot{\rho} = -i[\hat{H}(t), \rho] + \kappa\mathcal{D}[\hat{a}]\rho, \quad (10)$$

where $\mathcal{D}[\hat{x}]\rho = \hat{x}\rho\hat{x}^\dagger - \{\hat{x}^\dagger\hat{x}, \rho\}/2$ is the usual dissipator that describes evolution due to interaction with the environment; in this case, decay of oscillator photons into the output port at a rate κ . By following this approach, our simulated $s(t)$ is generated by a model for open-system evolution that accounts for the effects of input/output on the reservoir, including measurement-induced back action. Further, considering open-system evolution is useful as it can reduce susceptibility to overfitting, thereby improving performance.

As a model of the classical reservoir, we use the equation of motion for a single nonlinear oscillator given by

$$\dot{a} = -iK(1 + 2|a|^2)a - \frac{\kappa}{2}a - i\alpha u(t), \quad (11)$$

where $a = (X + iP)/\sqrt{2}$ is the complex scalar amplitude of the classical oscillator, with a^* its complex conjugate. Note that for the quantum model of Eq. (10), due to the Kerr nonlinearity and the noncommuting nature of the \hat{X} and \hat{P} operators, there are many classical models that can be derived from it. We have found that the choice of the specific classical model does not qualitatively affect the reservoirs' performances, and the classical model we have chosen gives good dynamical agreement with the quantum model in parameter regimes where this is expected ($K \ll \alpha, \kappa$).

IV. SINE WAVE PHASE ESTIMATION

The task we consider for our reservoirs is the estimation of the phase of an oscillatory signal. This fundamental task in signal processing is challenging as it embodies a nonlinear, nonconvex optimization problem [29]. Explicitly, the input to the reservoir is

$$u(t) = \sin(\omega_u t + \phi), \quad (12)$$

for a fixed, known frequency ω_u and an unknown phase ϕ in the interval $[0, \pi/2]$ [30]. This input signal corresponds to a double-sideband modulation of a single frequency carrier tone centered at the nonlinear oscillator resonance frequency Ω , as shown in Fig. 1. In practice, such an input can be implemented in many modalities, e.g., at microwave frequencies it can be achieved using a microwave mixer, as depicted in Fig. 1.

The task is to estimate the numerical value of this unknown phase using the reservoir computer. As such, the output weight “matrix” is a $1 \times N_{\text{out}}P$ vector, and the reservoir phase estimate is given by

$$\phi^{\text{est}} = \mathbf{W}_{\text{out}}\vec{s}_t. \quad (13)$$

The optimal weight matrix is determined by training on a set of known phases, and throughout this work we use a set of M equidistantly spaced phases in the interval $[0, \pi/2]$ as our training set.

We measure the performance of the reservoir at the sine wave phase estimation task via the root mean square (RMS) error of the estimated phases for a test set of size T , given by

$$r = \sqrt{\frac{1}{T} \sum_{j=0}^T |\phi_j^{\text{est}} - \phi_j^{\text{act}}|^2}, \quad (14)$$

where ϕ_j^{act} is the actual value of the phase for the j th element of the test set. The test set consists of uniformly random phases from the interval $[0, \pi/2]$.

V. RESULTS

We compare the performance of quantum and classical reservoirs with Gaussian distributed parameters K, κ, α , and ω_u , with average values $\bar{K}/\bar{\kappa} = -2$, $\bar{\omega}_u/\bar{\kappa} = 10$, $\bar{\alpha}/\bar{\kappa} = 6$, and $\bar{\kappa} = 1$ in arbitrary units. Each parameter's Gaussian distribution has a standard deviation that is 10% of its average value. Simulation of the reservoir evolution, Eqs. (10) and (11), is done in *Julia* [31], and the quantum reservoir simulation uses the master equation solver package *MESolve.jl* [32]. Each input signal lasts for a time duration of $2/\bar{\kappa}$, and we measure the output at a uniformly distributed set of $P = 100$ points across this time, i.e., an output sample rate of $50\bar{\kappa}$.

For the classical reservoirs, the number of computational nodes is $N = 2$, corresponding to the two quadratures X and P , or alternatively the real and imaginary parts of the complex amplitude a . For the quantum reservoirs, the maximum number of computational nodes scales as the Hilbert space dimension. While formally infinite for an oscillator, for our simulations we must truncate this to some finite value d , and a density matrix on such a space has $d^2 - 1$ real parameters. Thus, $N \leq d^2 - 1$, since we are not guaranteed that all E_{nm} will be independent of one another for the given quantum state. In both cases, we have $N_{\text{out}} = 1$ since we only use the expectation value of the X quadrature for output.

For all the results presented in the following, training is done on an equidistantly spaced grid of M training phases on the interval $[0, \pi/2]$. We find this improves performance compared to training with uniformly random phases from $[0, \pi/2]$, as the equidistant grid ensures that the reservoir is trained with phases that span the full interval and are representative of all possible inputs. Further, we do not fix the ridge regression parameter of Eq. (3) but find an optimal value in each unique case (quantum vs classical, training size M , Hilbert space dimension, etc.) by searching over values from $\gamma = 10^{-12}$ to $\gamma = 1$ in powers of ten, and report the smallest value of the RMS error found.

A. Quantum-classical improvement

We first compare the performance of the quantum and classical reservoirs as a function of training set size, as shown in Fig. 2, with a Hilbert space dimension of $d = 12$ used for the QRC simulations. For smaller training set sizes, especially at intermediate scale around 30, there is a multiple order-of-magnitude smaller average RMS error for the quantum reservoirs compared to the classical reservoirs. This is the first example of an improvement for QRC compared to the equivalent CRC model. This improvement exists for specific QRC/CRC parameter sets, and is true on average, but is not a universal property for all parameter values. There are parameter sets $(K, \omega, \alpha, \kappa)$ where CRC outperforms QRC, as is shown in Fig. 2 by the best case performance curves for CRC and QRC. Each point on these curves is the lowest RMS error for the RC model, and this best case performance does not occur for the same parameters at each training set size.

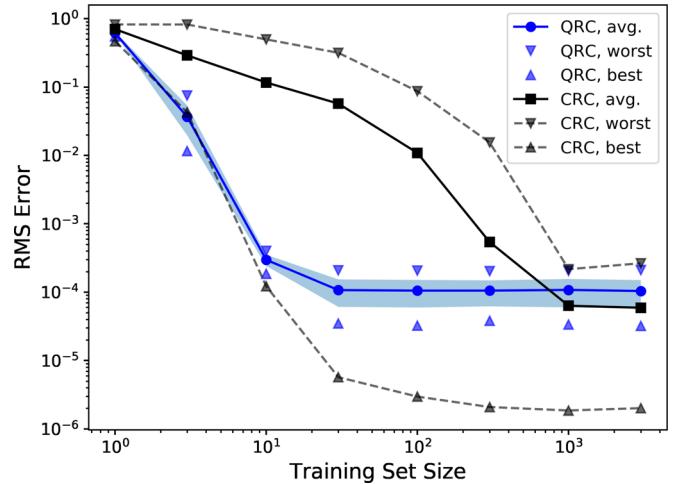


FIG. 2. Quantum and classical reservoir performance. As a function of training set size, average root mean square (RMS) error on the test set for the quantum nonlinear oscillator reservoirs (QRC) with Hilbert space dimension $d = 12$, and the classical nonlinear oscillator reservoirs (CRC). The average results are for 101 reservoirs with Gaussian distributed parameters around the mean values described in the main text. The reservoir output is sampled at 100 equally spaced points in time, and the test set size is 5000. One standard deviation for the QRC performance is shown in the shaded blue region. The RMS error for the QRC and CRC with best and worst case performance (determined independently for each training set size) is also shown.

Nevertheless, there is a demonstrable improvement for QRC over CRC for many specific parameter sets and on average, even if there is no improvement for QRC over all CRC in nonlinear oscillators.

The second important point highlighted in Fig. 2 is the performance spread of the QRC and CRC models, which we quantify by the RMS error spread factor

$$\text{RMSE Spread Factor} = \frac{\text{Best case} - \text{Worst case}}{\text{Average}}. \quad (15)$$

The QRC spread is at worst less than a factor of 2, as highlighted by the narrow shadowed region representing one standard deviation in RMS error. By comparison, the CRC spread is typically around a factor of 10 and can be as much as a factor of 30.

This demonstrates the second improvement of QRC over CRC nonlinear oscillators at sine phase estimation: reliability. As we have shown, the performance of the quantum reservoirs is not highly dependent on the specific oscillator parameters, while the performance of a given classical reservoir cannot be inferred from the performance of a different classical reservoir, even one with very similar oscillator parameters. This is important in simulation and design, where the QRC requires far less parameter optimization to find an effective reservoir [33–35], as well as in practical implementation, which will naturally have a spread in parameters due to errors in device fabrication, control, and measurement.

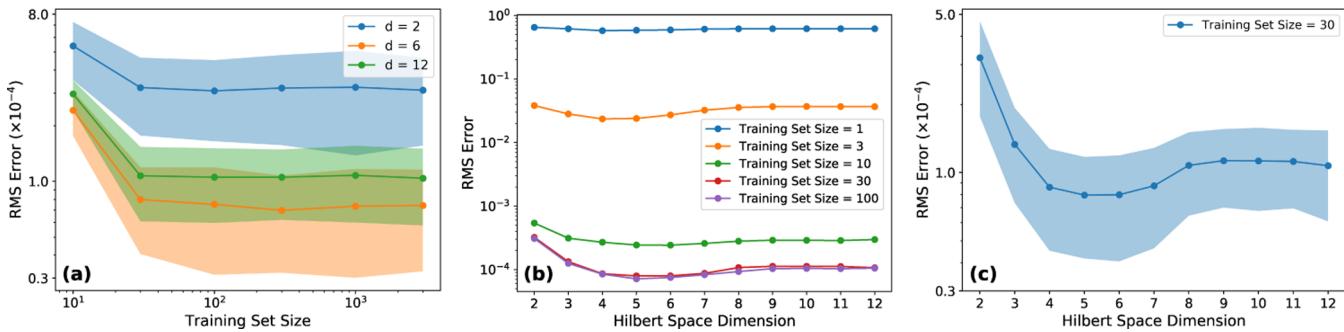


FIG. 3. Hilbert space dimension effect on QRC performance. (a) Average RMS error as a function of training set size for quantum nonlinear oscillator reservoirs with various Hilbert space dimension. One standard deviation for 101 reservoirs with Gaussian distributed parameters is shown in the shaded regions. (b),(c) Average RMS error as a function of Hilbert space dimension, for (b) various training set sizes and (c) for a training set size of 30, which is near optimal. In (c) one standard deviation is shown in the shaded area.

B. Hilbert space dimension dependence

To elucidate the source of the quantum-classical improvements, we begin by studying how the performance of the quantum reservoirs depends on the Hilbert space dimension of the simulations. The results shown in the previous subsection are for a dimension chosen such that the simulations are a good approximation to the infinite dimensional oscillator. This is quantified by the fact that for the quantum state $\rho(t)$ in the oscillator during our simulations, the higher Fock states are minimally occupied, and the canonical commutation relation $\text{Tr}([\hat{a}, \hat{a}^\dagger]\rho(t)) = 1$ is valid with less than 1% error at all times. As we decrease Hilbert space dimension in our study, the quantum model stops being a valid description of a nonlinear oscillator but remains a valid description of a qudit, of which many physical examples exist [36,37].

Figure 3 shows results for QRC models ranging from a qubit, $d = 2$, to the approximate nonlinear oscillator of the previous subsection, $d = 12$. In all cases, we continue to use measurements of the \hat{X} operator as the output, which in the qubit case becomes the Pauli operator $\hat{\sigma}_x$. As Fig. 3(a) shows, the average RMS error depends on Hilbert space dimension in a nonmonotonic fashion, while the spread remains roughly consistent. This is further demonstrated in Fig. 3(b), with the RMS error showing a local minimum around $d = 5$ for all training set sizes. Focusing on a training set size of 30 in Fig. 3(c), we see that the best performance occurs at $d = 5$ and $d = 6$, which is almost a factor of 4 better than the worst performance at $d = 2$.

The increase in performance as Hilbert space dimension increases is to be expected in QRC, as an increase in Hilbert space dimension implies an increase in the number of degrees of freedom of the reservoir, and thus, its computational power. The nonlinearity of the system dynamics is key in this regard, as it can create non-Gaussian states in the oscillator. Only for such states can the expectation values E_{nm} of Eq. (4) for $n, m > 2$ be independent of lower order E_{nm} and thus expand the computational space of the reservoir.

The existence of a local minimum, and the saturation of performance for $d \geq 8$, is likely explained by a combination of effects. Firstly, due to the finite decay rate κ , the expectation values E_{nm} , or equivalently the Fock state populations, decay exponentially with a rate that is linearly proportional to photon

number. As such, while higher Hilbert space dimension offers more computational variables, information is lost from them at a faster rate.

Secondly, even without finite κ , information can be lost in the higher order E_{nm} by an information scrambling argument. Since we only measure $E_{10} = \langle \hat{X} \rangle$, it is possible that information relevant to sine phase estimation is spread into other E_{nm} and is not accessible in E_{10} during the finite times at which we measure. Finally, the ratio of the drive amplitude α to the drive frequency ω determines how much energy can enter the oscillator, and as higher Fock states are occupied, the nonlinearity K plays a role as well. This further restricts how much of the computational space of the reservoir is accessible.

While we have good heuristic explanations for the weak dependence of QRC performance on Hilbert space dimension, these cannot explain the significant quantum-classical improvement in both average performance and reliability for intermediate training set sizes seen in Fig. 2. These improvements persist even when the QRC is a qubit. This is particularly intriguing as a single qubit is a *classical* system, as discussed below, and we would not have expected any improvement in this case. In the following subsection we delve into this surprising result further and show how it highlights the likely explanation for the majority of the quantum-classical improvement seen for any Hilbert space dimension.

C. Origin of the quantum-classical improvement

When we say that a single qubit is a classical system, we mean that there is a local hidden variable model that can completely describe the state space and dynamical evolution of a single qubit. For our purposes, we will consider the Kochen-Specker model [38], which maps a density matrix ρ for a single qubit to the three classical variables (r, θ, ϕ) that parametrize a unit sphere in spherical coordinates. We consider this hidden variable model for a qubit as a reservoir (HVRC) and simulate its performance at the sine phase estimation task. The dynamical evolution of the HVRC is described by a set of transcendental differential equations, so for simplicity we use the results of the quantum simulations for a qubit and convert to the HV model parameters using the

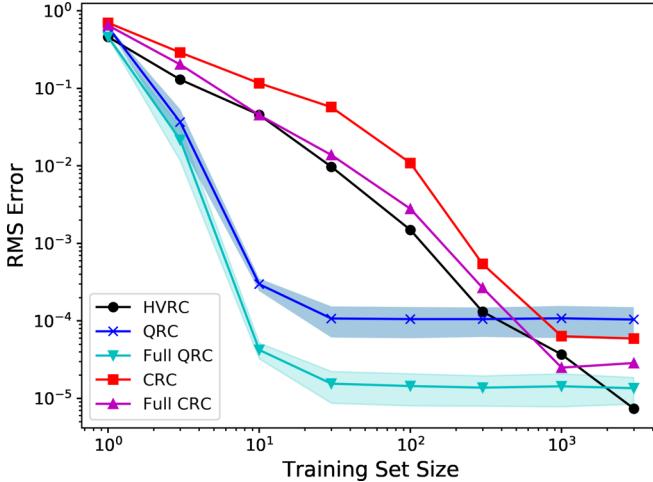


FIG. 4. Reservoir performance for qubit and classical models. As a function of training set size, average root mean square (RMS) error on the test set for a qubit reservoir, either with a single measured expectation value (QRC) or full tomography (full QRC), the hidden variable model for a qubit as a reservoir (HVRC), and classical nonlinear oscillator reservoirs with single quadrature (CRC) or both quadrature (full CRC) measurements. Reservoir parameter set, measurement sampling, and test set size are the same as in Fig. 2. One standard deviation for the QRC and full QRC performance is shown in the shaded regions.

relationships

$$(x, y, z) = (\text{Tr}[\hat{\sigma}_x \rho], \text{Tr}[\hat{\sigma}_y \rho], \text{Tr}[\hat{\sigma}_z \rho]),$$

$$(r, \theta, \phi) = \left(x^2 + y^2 + z^2, \arccos\left(\frac{z}{r}\right), \arctan\left(\frac{y}{x}\right) \right),$$

which one may recognize as the familiar coordinate transformation from Cartesian to spherical coordinates.

We simulate the performance of the HVRC under the same conditions (parameter sets, training and test sets) as for the QRC and CRC presented previously and use all three variables as the output. Figure 4 shows the average performance of the various RC models, and as can be seen, there is considerable performance improvement for the QRC compared to the HVRC at intermediate training set size. There is an even larger improvement for what we call full QRC, which uses full tomographic data on the qubit state (expectation values for all three Pauli operators) as the output, as opposed to QRC which uses only $\text{Tr}[\hat{\sigma}_x \rho]$.

We emphasize that all models considered in this subsection describe reservoir dynamics that are purely classical physics. The only difference between the QRC models and the HVRC model is that the HVRC model uses the output variables (r, θ, ϕ) , while the QRC models use the output variables $(x, y, z) = (r \sin \theta \cos \phi, r \sin \theta \sin \phi, r \cos \theta)$, which are a nonlinear transformation of the HVRC output. As such, the average performance improvement between QRC and HVRC can only be attributed to this nonlinear transformation. Such a nonlinear transformation cannot be implemented by the linear processing of output data that occurs in the typical operation of a reservoir computer, see Eq. (2). Nevertheless, previous work has shown that appending nonlinear post-processing to the output of a reservoir can have significant impact on

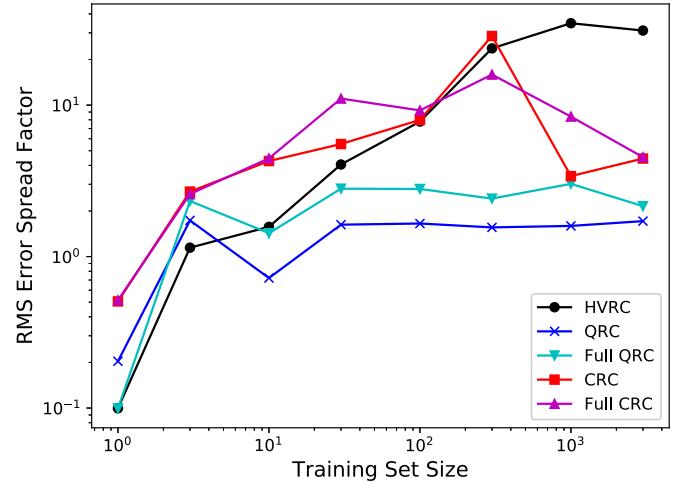


FIG. 5. Performance spread for qubit and classical models. As a function of training set size, spread in root mean square (RMS) error, cf. Eq. (15), on the test set for a qubit reservoir, either with a single measured expectation value (QRC) or full tomography (full QRC), the hidden variable model for a qubit as a reservoir (HVRC), and classical nonlinear oscillator reservoirs with single quadrature (CRC) or both quadrature (full CRC) measurements. Output sampling and test set size the same as in Figs. 2 and 4.

performance [39], though one may argue that this is costly to implement.

This nonlinear processing of data occurs intrinsically for the quantum reservoir due to quantum measurement and is therefore part of QRC itself, being responsible for the generation of output data. As in all RC, standard linear post-processing of the output still occurs, but we essentially get an otherwise costly nonlinear processing stage “for free” in QRC. We note that performance improvement due to measurement-induced nonlinearity has also been observed in classical hardware reservoirs, see, e.g., Refs. [24,40,41].

As we have shown, the nonlinearity of quantum measurement is the source of performance improvement between the QRC models and the HVRC for a qubit, from which we infer that it is also the main source of improvement of the qubit QRC compared to the CRC models (see Fig. 4, where full CRC uses both quadrature measurements as output). However, the average HVRC performance is better than that for CRC, which implies it is a better reservoir for the sine phase estimation task. Thus, part of the performance improvement of qubit QRC over CRC is also due to the fact that the HVRC underlying the QRC is better suited for the task at hand.

While we cannot make definitive statements for QRC beyond $d = 2$, as no local hidden variable models exist, the results for the qubit QRC indicate that the quantum-classical performance improvement seen in Fig. 2 is likely due to the nonlinearity of quantum measurement. Similarly, as shown in Fig. 5, which plots the RMS error spread factor of Eq. (15), we see that HVRC has a large spread in performance, similar to CRC. This indicates that the reliability of QRC (at all Hilbert space dimensions) also stems from the nonlinearity of quantum measurement. Thus, due to the improvement in both average performance and reliability, we believe the nonlinear oscillator QRC possesses an advantage over CRC due to the intrinsic nonlinearity of quantum measurement.

D. Input and output noise

As a final consideration for the performance of the nonlinear oscillator RC models, we examine the impact of noise in the both the output data and the input signal. Noise can have a variety of sources, such as thermal noise in the control and readout chains of the device, and will inevitably plague hardware implementations of our models. In particular, output noise is unavoidable for QRC due to the intrinsic uncertainty of quantum measurement. QRC uses the expectation value of an operator as its output variable, which, in a single-shot measurement, can only be estimated up to quantum uncertainty, often characterized by the standard deviation of the operator expectation value. For the \hat{X} operator we have chosen, this is given by

$$\Delta E_{\text{out}} = \sqrt{\text{Tr}[\rho \hat{X}^2] - \text{Tr}[\rho \hat{X}]^2}. \quad (16)$$

For a quantum system with a large (infinite) Hilbert space, such as a nonlinear oscillator, the system can be in a state (e.g., a coherent state) such that $|E_{01}| \gg |\Delta E_{01}|$. In this case, single-shot measurements give a good estimate of the output variable.

The situation is very different for low-dimensional systems. In the extreme case of a qubit, each measurement gives a binary output, and for most quantum states the standard deviation of a Pauli operator measurement is of the same order of magnitude as the expectation value. Thus, it is necessary to average over repeated runs of the same training/test case, to produce an estimate of the output variable \tilde{E}_{out} with the standard error of this estimate scaling as

$$\Delta \tilde{E}_{\text{out}} = \frac{\Delta E_{\text{out}}}{\sqrt{R}}, \quad (17)$$

where R is the number of repeated runs. The slow improvement with \sqrt{R} is unfavorable in practice, and importantly, is a feature of all quantum neuromorphic and machine learning protocols that use qubits and rely on real-valued (as opposed to binary) output data. Recently, it was pointed out that one must include the cost of these repetitions when assessing algorithm run time and scaling [17].

Returning to the nonlinear oscillator models, it should be clear that we use such models in part to reduce the necessity for repetition and keep R as small as possible. We artificially introduce output noise by adding a Gaussian random variable with zero mean and standard deviation σ to each sampled time point of the output from the reservoir. We do this for the training set, before calculating the weights using Eq. (3), and for the test set. Due to numerical simulation constraints, we consider a single set of reservoir parameters given by the mean values discussed at the beginning of Sec. V.

The dependence of QRC performance on output noise is shown in Fig. 6. As is to be expected, the RMS error increases monotonically as a function of the output noise. We measure the output noise in units of the input signal amplitude α , since the output signal expectation value will be a function of α . The results of Fig. 6 indicate that some repetition will be required for high performance in an experimental implementation of this nonlinear oscillator QRC, as it is unlikely that the combination of quantum measurement noise and other classical output noise sources can be brought below 10^{-3} (in units of the signal amplitude).

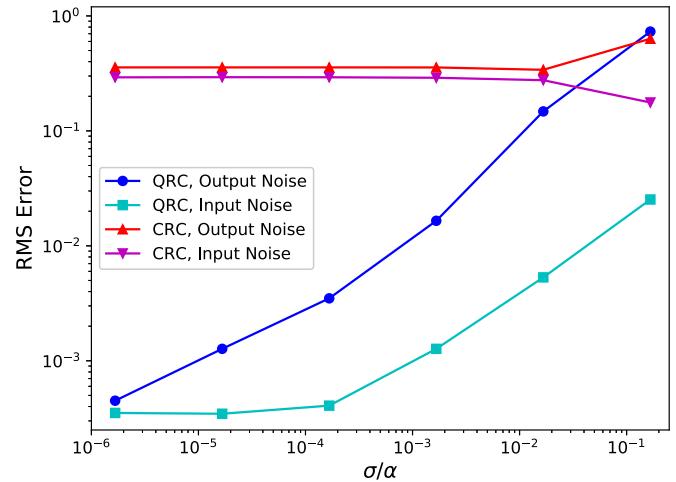


FIG. 6. Quantum and classical reservoir performance with noise. Root mean square (RMS) error on the test set for the quantum nonlinear oscillator reservoirs (QRC) with Hilbert space dimension $d = 12$, and the classical nonlinear oscillator reservoirs (CRC), as a function of noise standard deviation in units of the input drive amplitude (σ/α) for noise injected in either the output or the input. These results are for a single set of reservoir parameters described in the main text. Output sampling is as in Fig. 2, the training set size is 10, and the test set size is 5000 for output noise and 100 for input noise.

However, numerical simulation resources limit the maximum input signal amplitude we can consider for our reservoirs. It is likely that by increasing input signal amplitude, and therefore output signal amplitude, we can improve the performance of QRC in the presence of output noise, as the output noise would become a diminishing fraction of the output signal. An alternative approach would be to consider nonlinear oscillators with *in situ* parametric driving, such that they act as both a self-amplifier of their output signal and reduce quantum measurement noise by squeezing. The study of both these routes towards output noise insensitivity in QRC will be the subject of future theoretical and experimental effort.

Now turning to input noise, we artificially inject noise into the input signal to our quantum reservoir. To ensure that our simulations observe causality, we cannot use the high-performance adaptive time-step differential equation solvers that are default in `MESolve.jl`, but must switch to a considerably slower fixed time-step solver. We use a time-step $\delta t \ll 1/(50\bar{\kappa})$, $1/\bar{\omega}$ that is much smaller than the output sample time step, the oscillator lifetime, and the input signal oscillation period to ensure the noise closely mimics white noise with no temporal correlations. We add independent Gaussian random noise with zero mean and standard deviation $\sigma = s\sqrt{\delta t}$ at each time step, such that our simulations closely approximate the effects of Gaussian white noise in the input signal, where s is a scale factor describing the noise spectral amplitude.

Figure 6 also shows the results of our input noise simulations, which as can be seen, are much more favorable than for output noise. In particular, we do not observe an increase in RMS error until σ/α reaches a level comparable to the RMS error with no noise, after which the RMS error increases

roughly linearly. This indicates that the internal reservoir dynamics do not appear to amplify the effect of input noise and that training is less disrupted by small amounts of input noise compared to output noise. A detailed understanding of the mechanisms behind these observations is the subject of future study, but we do not expect input noise at experimentally accessible levels to have a limiting impact on QRC for our chosen task.

VI. CONCLUSION

In this work, we have introduced an approach to quantum reservoir computing that uses a single nonlinear oscillator as the reservoir. The computational nodes are formed by the complex amplitudes of the eigendecomposition of the system state, or equivalently, the expectation values of a complete basis of observables. We have demonstrated that this quantum reservoir has improved performance compared to the equivalent classical one at sine phase estimation, both in terms of average estimation error and in reliability of performance across reservoirs with different internal parameters. Sine phase estimation is a foundational task in signal processing and a pragmatic example of an estimation task. Given the performance demonstrated by our QRC model at this task, it is likely that it or a similar QRC model will perform well at other signal processing or estimation tasks.

By studying the performance dependence on Hilbert space dimension down to the single qubit level, and comparing to the fully-classical hidden variable theory describing the dynamical evolution of a qubit, we have determined that the main source of the quantum-classical improvement is the nonlinearity inherent to quantum measurement. The impact of this extends beyond reservoir computing, and to our knowledge this has not previously been identified as a beneficial feature intrinsic to all quantum neuromorphic computing approaches. We believe this further divides quantum and classical neuromorphic and machine learning methodologies, strengthens the case for potential (heuristic) quantum advantage, and may shed new light on previous results [42].

With an eye towards experimental implementation, we studied the impact of injected input or output noise on performance and have found that output noise is the more detrimental. While we have proposed several mitigation strategies, it appears that some repetition of experiments will be necessary to obtain accurate expectation values, as is the case in most proposed implementations of quantum neuromorphic computing or quantum machine learning. Nevertheless, we expect the required number of repetitions to scale more favorably for our continuous variable reservoir than for discrete variable approaches.

The nonlinear oscillator reservoir we have proposed is well suited to an implementation in realistic experiments. In particular, the tools and techniques of circuit quantum electrodynamics (cQED) [43] that have been developed to build conventional quantum processors provide an attractive realization of the building blocks needed for QRC. cQED experiments revolve around the control of the quantum degrees of freedom of well-isolated nonlinear oscillators (e.g., superconducting qubits), and their basic ingredient, the Josephson junction, affords a high degree of control over circuit parameters [37,44], as well as strong intrinsic nonlinearity. Further, both nonlinear opto- or electromechanical resonators [45–47] and nonlinear photonic systems [48,49] are attractive test beds for continuous variable QRC. We expect that experimental implementations of the ideas presented in this work will provide a promising area of applications for devices heretofore used in quantum information research.

ACKNOWLEDGMENTS

The authors acknowledge insightful discussions on reservoir computing with Daniel Gauthier. This material is based upon work supported by the US Army Research Office under Contract No: W911NF-19-C-0092. Any opinions, findings and conclusions or recommendations expressed in this material are those of the authors and do not necessarily reflect the views of the US Army Research Office.

- [1] M. A. Nielsen and I. L. Chuang, *Quantum Computation and Quantum Information: 10th Anniversary Edition* (Cambridge University Press, NY, 2010).
- [2] F. Arute *et al.*, Quantum supremacy using a programmable superconducting processor, *Nature (London)* **574**, 505 (2019).
- [3] B. M. Terhal, Quantum error correction for quantum memories, *Rev. Mod. Phys.* **87**, 307 (2015).
- [4] E. T. Campbell, B. M. Terhal, and C. Vuillot, Roads towards fault-tolerant universal quantum computation, *Nature (London)* **549**, 172 (2017).
- [5] J. Preskill, Quantum Computing in the NISQ era and beyond, *Quantum* **2**, 79 (2018).
- [6] D. Ristè, M. P. da Silva, C. A. Ryan, A. W. Cross, A. D. Córcoles, J. A. Smolin, J. M. Gambetta, J. M. Chow, and B. R. Johnson, Demonstration of quantum advantage in machine learning, *npj Quantum Inf.* **3**, 16 (2017).
- [7] C. D. Schuman, T. E. Potok, R. M. Patton, J. D. Birdwell, M. E. Dean, G. S. Rose, and J. S. Plank, A survey of neuromorphic computing and neural networks in hardware, [arXiv:1705.06963](https://arxiv.org/abs/1705.06963).
- [8] W. Maass, T. Natschläger, and H. Markram, Real-time computing without stable states: A new framework for neural computation based on perturbations, *Neural Comput.* **14**, 2531 (2002).
- [9] H. Jaeger and H. Haas, Harnessing nonlinearity: Predicting chaotic systems and saving energy in wireless communication, *Science* **304**, 78 (2004).
- [10] D. Verstraeten, B. Schrauwen, M. D'Haene, and D. Stroobandt, An experimental unification of reservoir computing methods, *Neural Networks* **20**, 391 (2007).
- [11] K. Fujii and K. Nakajima, Harnessing Disordered-Ensemble Quantum Dynamics for Machine Learning, *Phys. Rev. Appl.* **8**, 024030 (2017).
- [12] S. Ghosh, A. Opala, M. Matuszewski, T. Paterek, and T. C. H. Liew, Quantum reservoir processing, *npj Quantum Inf.* **5**, 35 (2019).
- [13] A. Kutvonen, K. Fujii, and T. Sagawa, Optimizing a quantum reservoir computer for time series prediction, *Sci. Rep.* **10**, 14687 (2020).

[14] M. Schuld and N. Killoran, Quantum Machine Learning in Feature Hilbert Spaces, *Phys. Rev. Lett.* **122**, 040504 (2019).

[15] J. Chen and H. I. Nurdin, Learning nonlinear input–output maps with dissipative quantum systems, *Quant. Info. Proc.* **18**, 198 (2019).

[16] K. Nakajima, K. Fujii, M. Negoro, K. Mitarai, and M. Kitagawa, Boosting Computational Power through Spatial Multiplexing in Quantum Reservoir Computing, *Phys. Rev. Appl.* **11**, 034021 (2019).

[17] L. G. Wright and P. L. McMahon, The capacity of quantum neural networks, in *Conference on Lasers and Electro-Optics*, OSA Technical Digest (Optical Society of America, 2020), paper JM4G.5.

[18] S. Ghosh, T. Paterek, and T. C. H. Liew, Quantum Neuromorphic Platform for Quantum State Preparation, *Phys. Rev. Lett.* **123**, 260404 (2019).

[19] S. Ghosh, T. Krisnanda, T. Paterek, and T. C. H. Liew, Universal quantum reservoir computing, [arXiv:2003.09569](https://arxiv.org/abs/2003.09569).

[20] M. Negoro, K. Mitarai, K. Fujii, K. Nakajima, and M. Kitagawa, Machine learning with controllable quantum dynamics of a nuclear spin ensemble in a solid, [arXiv:1806.10910](https://arxiv.org/abs/1806.10910).

[21] J. Chen, H. I. Nurdin, and N. Yamamoto, Temporal Information Processing on Noisy Quantum Computers, *Phys. Rev. Appl.* **14**, 024065 (2020).

[22] S. Dasgupta, K. E. Hamilton, P. Lougovski, and A. Banerjee, Designing a nisq reservoir with maximal memory capacity for volatility forecasting, [arXiv:2004.08240](https://arxiv.org/abs/2004.08240).

[23] G. Dion, S. Mejaouri, and J. Sylvestre, Reservoir computing with a single delay-coupled non-linear mechanical oscillator, *J. Appl. Phys.* **124**, 152132 (2018).

[24] L. Appeltant, M. C. Soriano, G. Van der Sande, J. Danckaert, S. Massar, J. Dambre, B. Schrauwen, C. R. Mirasso, and I. Fischer, Information processing using a single dynamical node as complex system, *Nat. Commun.* **2**, 468 (2011).

[25] 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, *Phys. Rev. E* **91**, 020801(R) (2015).

[26] J. Torrejon, M. Riou, F. A. Araujo, S. Tsunegi, G. Khalsa, D. Querlioz, P. Bortolotti, V. Cros, K. Yakushiji, A. Fukushima, H. Kubota, S. Yuasa, M. D. Stiles, and J. Grollier, Neuromorphic computing with nanoscale spintronic oscillators, *Nature (London)* **547**, 428 (2017).

[27] C. W. Gardiner and P. Zoller, *Quantum Noise* (Springer, Heidelberg, 2004).

[28] H. P. Breuer and F. Petruccione, *The Theory of Open Quantum Systems* (Oxford University Press, New York, 2002).

[29] S. M. Kay, *Fundamentals of Statistical Signal Processing. v.1: Estimation Theory* (Prentice Hall, NJ, 1993).

[30] We find that performance decreases significantly if we increase the interval size. We attribute this to the reservoir confusing phases in $[0, \pi/2]$ with those in $[\pi/2, \pi]$ for which $|\sin(\phi)|$ is the same. This may be related to the parity of the reservoir nonlinearity and is a topic of future study.

[31] J. Bezanson, A. Edelman, S. Karpinski, and V. B. Shah, Julia: A fresh approach to numerical computing, *SIAM Rev.* **59**, 65 (2017).

[32] Luke C. G. Govia (2020, April 7), BBN-Q/MESolve.jl: v1.0.0 release (Version v1.0.0), Zenodo, <http://doi.org/10.5281/zenodo.3743496>.

[33] H. Jaeger, M. Lukoševičius, D. Popovici, and U. Siewert, Optimization and applications of echo state networks with leaky-integrator neurons, *Neural Networks* **20**, 335 (2007).

[34] J. Yperman and T. Becker, Bayesian optimization of hyperparameters in reservoir computing, [arXiv:1611.05193](https://arxiv.org/abs/1611.05193).

[35] A. Griffith, A. Pomerance, and D. J. Gauthier, Forecasting chaotic systems with very low connectivity reservoir computers, *Chaos: Interdisciplinary J Nonlin. Sci.* **29**, 123108 (2019).

[36] J. Clarke and F. K. Wilhelm, Superconducting quantum bits, *Nature (London)* **453**, 1031 (2008).

[37] P. Krantz, M. Kjaergaard, F. Yan, T. P. Orlando, S. Gustavsson, and W. D. Oliver, A quantum engineer’s guide to superconducting qubits, *Appl. Phys. Rev.* **6**, 021318 (2019).

[38] S. Kochen and E. P. Specker, The problem of hidden variables in quantum mechanics, *J. Math. Mech.* **17**, 59 (1967).

[39] F. Abreu Araujo, M. Riou, J. Torrejon, S. Tsunegi, D. Querlioz, K. Yakushiji, A. Fukushima, H. Kubota, S. Yuasa, M. D. Stiles, and J. Grollier, Role of non-linear data processing on speech recognition task in the framework of reservoir computing, *Sci. Rep.* **10**, 328 (2020).

[40] Y. Paquot, F. Duport, A. Smerieri, J. Dambre, B. Schrauwen, M. Haelterman, and S. Massar, Optoelectronic reservoir computing, *Sci. Rep.* **2**, 287 (2012).

[41] D. Brunner, M. C. Soriano, C. R. Mirasso, and I. Fischer, Parallel photonic information processing at gigabyte per second data rates using transient states, *Nat. Commun.* **4**, 1364 (2013).

[42] C. M. Wilson, J. S. Otterbach, N. Tezak, R. S. Smith, A. M. Polloreno, P. J. Karalekas, S. Heidel, M. S. Alam, G. E. Crooks, and M. P. da Silva, Quantum kitchen sinks: An algorithm for machine learning on near-term quantum computers, [arXiv:1806.08321](https://arxiv.org/abs/1806.08321).

[43] A. Blais, J. Gambetta, A. Wallraff, D. I. Schuster, S. M. Girvin, M. H. Devoret, and R. J. Schoelkopf, Quantum-information processing with circuit quantum electrodynamics, *Phys. Rev. A* **75**, 032329 (2007).

[44] 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, *Phys. Rev. A* **76**, 042319 (2007).

[45] J. C. Sankey, C. Yang, B. M. Zwickl, A. M. Jayich, and J. G. E. Harris, Strong and tunable nonlinear optomechanical coupling in a low-loss system, *Nat. Phys.* **6**, 707 (2010).

[46] S. Rips, I. Wilson-Rae, and M. J. Hartmann, Nonlinear nanomechanical resonators for quantum optoelectromechanics, *Phys. Rev. A* **89**, 013854 (2014).

[47] M.-A. Lemonde, N. Didier, and A. A. Clerk, Enhanced nonlinear interactions in quantum optomechanics via mechanical amplification, *Nat. Commun.* **7**, 11338 (2016).

[48] D. V. Strekalov, C. Marquardt, A. B. Matsko, H. G. L. Schwefel, and G. Leuchs, Nonlinear and quantum optics with whispering gallery resonators, *J. Opt.* **18**, 123002 (2016).

[49] G. Lin, A. Coillet, and Y. K. Chembo, Nonlinear photonics with high-q whispering-gallery-mode resonators, *Adv. Opt. Photon.* **9**, 828 (2017).