PHILOSOPHICAL TRANSACTIONS A

royalsocietypublishing.org/journal/rsta

Research

Cite this article: Gyurik C, Wudarski F, Philip EJ, Sannia A, Sadeghi H, Kyriienko O, Venturelli D, Gentile AA. 2026 From quantum feature maps to quantum reservoir computing: an applicative perspective. *Phil. Trans. R. Soc. A* **384**: 20250085.
https://doi.org/10.1098/rsta.2025.0085

Received: 8 June 2025
Accepted: 23 October 2025

One contribution of 9 to a discussion meeting issue 'Bits, neurons and qubits for sustainable AI'.

Subject Areas:
quantum physics, quantum computing

Keywords:
reservoir computing, quantum computing, neutral atoms

Author for correspondence:
Antonio Andrea Gentile
e-mail: andrea.gentile@pasqal.com

†These authors contributed equally to the study.

THE ROYAL SOCIETY PUBLISHING

From quantum feature maps to quantum reservoir computing: an applicative perspective

Casper Gyurik[1,†], Filip Wudarski[2,†], Evan John Philip[1], Antonio Sannia[2,3], Hossein Sadeghi[4], Oleksandr Kyriienko[5], Davide Venturelli[2] and Antonio Andrea Gentile[1]

[1]Pasqal SaS, Amsterdam, The Netherlands
[2]USRA, Washington, DC, USA
[3]UIB-CSIC, Palma de Mallorca, Spain
[4]Pasqal SaS, Boston, MA, USA
[5]The University of Sheffield Department of Physics and Astronomy, Sheffield, UK

FW, 0000-0002-0911-8342; EJP, 0000-0001-5405-2306; AAG, 0000-0002-1763-9746

We explore the interplay between two emerging paradigms: reservoir computing (RC) and quantum computing (QC). We observe how quantum systems featuring beyond-classical correlations and vast computational spaces can serve as non-trivial, experimentally viable reservoirs for typical tasks in machine learning. With a focus on neutral-atom quantum processing units, we describe and exemplify a novel quantum RC (QRC) workflow. We conclude by exploratively discussing the main challenges ahead, while arguing how QRC can offer a natural candidate to push forward RC applications.

This article is part of the discussion meeting issue 'Bits, neurons and qubits for sustainable AI'.

1. Introduction

Quantum computing (QC) is a specific model of computation that uses the principles of quantum mechanics to process information. The fundamental unit to store and process information in such a model is the *qubit*, whose

isolated state can be fully described by two real numbers unlike the binary information carried by single bits.

Quantum systems exhibit interference and entanglement, which allow qubits to be correlated in ways that classical bits cannot. Digital computation in this model is carried out using quantum gates, which are operations that change the state of qubits, similar to how logic gates operate on bits in classical circuits. Analogue operations where the qubits are evolved according to their intrinsic system behaviour are also possible in quantum circuits.

While quantum computers grow in size with a number of qubits reaching hundreds, they are still sparsely connected and have non-negligible error rates [1]. Hybrid quantum–classical approaches offer the possibility of enhancing the capabilities of quantum devices and tailoring them for specific applications [2,3]. These algorithms show modularity and flexibility, allowing them to exploit existing hardware while delivering the best performance. In addition, one can scale the size of the system in a modular way, connecting registers with a handful of qubits into correlated multi-register systems with hundreds. The former can serve as a test bed for theoretical investigations, as well as numerical simulations, that could be used as a performance benchmark for quantum hardware runs. The latter regime, instead, is widely believed to include instances that cannot be simulated classically, leading to the concept of *quantum advantage* [4].

QC can be seen as a part of a broader field of physics-based computing, also encompassing analogue, reservoir and neuromorphic frameworks. In particular, reservoir computing (RC) leverages a complex dynamical system to map inputs into a high-dimensional computational space and processes readout data to efficiently solve learning tasks, including time-series predictions and classification [5,6]. Neuromorphic computing explores novel ways to process information, attempting to replicate brain processes with the intent of matching its efficiency. In this regard, some routes towards neuromorphic computing entertain the adoption of reservoir-like mappings enabled by biological or physical systems [7]. RC approaches exemplify heuristic computational paradigms. These methods typically lack formal performance guarantees but rely on empirical validation, offering solutions to complex or high-dimensional tasks that otherwise are prohibitively expensive [8].

The main appeal of RC is twofold: (i) because only the linear readout is trained, the models are relatively easy to optimize while still potentially capturing temporal dependencies (unlike recurrent neural networks, which are notoriously difficult to train owing to challenges in backpropagation); and (ii) with appropriately chosen reservoirs, both the computational and energy costs of training and inference can be much lower compared to other approaches [6,8]. The promise of QC resides in the favourable scaling of computational resources, compared to the best-known classical methods, expected when quantum algorithms are applied to certain problems, such as factoring large numbers or simulating quantum systems, making QC an obvious candidate in the efforts to optimize the computing load across several disciplines [4].

In this contribution, we explore quantum RC (QRC) as an emerging machine learning paradigm that combines quantum evolution with non-trivial, high-dimensional feature processing. We concentrate on physical systems that are based on neutral-atom arrays, with arbitrary connectivity and strong interactions. Such systems can naturally simulate complex multi-qubit models and their dynamics, satisfying one of the main requirements of RC. We describe in full generality the QRC paradigm and highlight the differences from other quantum machine learning (QML) approaches. We complement such analysis with exemplary applications of neutral atoms to relevant tasks in QRC, suggesting a novel hybrid scheme to leverage noisy intermediate-scale quantum (NISQ) devices for early implementations.

Building on such discussion, we argue how QC can amplify other non-von Neumann models of computation, providing benefits from coherent operation and entanglement, and in particular offer a promising venue to expand efforts in the direction of reservoir and neuromorphic computing paradigms.

2. Basics of reservoir computing

RC systems have long been proposed and successfully applied across a wide range of domains. As outlined in [6], these applications can be broadly categorized into a few main areas: signal classification (e.g. spoken digit recognition), time-series prediction (e.g. stock market forecasting), control of dynamical systems (e.g. real-time robot control) and partial differential equation-based computations (e.g. fast simulation of certain physical systems governed by differential equations).

Formally, an RC system operates as follows. Let the input data be denoted by $u_1, \ldots, u_t$, where each $u_i \in \mathbb{R}^n$, and let $u_i(j)$ denote the jth entry of u_i. We index the data according to the order in which it is presented to the learner, to emphasize potential temporal dependencies. The dynamical map that describes the reservoir evolution is denoted by $\mathcal{R}$, and in typical digital implementations, it is parameterized by a set of (fixed) random matrices W. At each time step t, the reservoir state r_t is given by the recursive relation

$$r_t = \mathcal{R}(u_t, r_{t-1}; W) \in \mathbb{R}^N. \tag{2.1}$$

For a suitable choice of $\mathcal{R}$, r_t is a function of the past inputs up to a typical characteristic time scale, t_{FM}, known as fading memory time, e.g.

$$r_t = r_t(\{u_i\}_{t-t_{\mathrm{FM}}}^{t}), \tag{2.2}$$

allowing the system to encode temporal dependencies in its state. The sharp cut-off invoked in equation (2.2) is often a reasonable approximation to more physically sound fading mechanisms, and hence we adopt it for the purposes of our work. Alternative RC investigations, however, include external, potentially time-unbounded working memories [9].

Finally, a linear readout layer $W_{\mathrm{out}} \in \mathbb{R}^{m \times N}$ is trained to produce

$$y_t = W_{\mathrm{out}} r_t, \tag{2.3}$$

where $y_t \in \mathbb{R}^m$ is the output, i.e. the reservoir prediction at time t. Therefore, we can distinguish three main subroutines: (i) data embedding, (ii) dynamical mapping and (iii) post-processing for training and prediction. When studying what makes a good reservoir, there are certain key considerations.

A crucial one is whether the reservoir exhibits some form of *fading memory*. This requirement is twofold: first, the state of the reservoir should depend on past inputs. However, the reservoir should primarily depend on recent inputs rather than inputs from the distant past. The first aspect is typically quantified by the *short-term memory capacity* of the system [10], while the second is referred to as the *echo state property* [11] and can be measured, for instance, using the echo state property index [12]. The fading memory is imprinted naturally in the dynamical map equation (2.2), which exploits nonlinear activation functions crucial for good performance.

Equally important is the presence of *nonlinearity* in the reservoir internal dynamics. Since the readout layer is typically linear, like in equation (2.3), the reservoir must introduce sufficient complexity to allow for non-trivial transformations of the input. This often requires tuning the system close to the edge of chaos—a regime where the dynamics is rich and sensitive to inputs, but not so unstable that it becomes unpredictable, or input information is completely scrambled. This balance ensures a high-dimensional, expressive training space while keeping enough structure for meaningful processing.

A significant portion of research in RC is dedicated to identifying physical systems that can effectively implement the reservoir mapping r specified in equation (2.1). As detailed in [6], various platforms have been explored for this purpose, including but not limited to: memristors [13], spintronic devices [14], photonic systems [15], digital field programmable gate arrays and application-specific integrated circuits [16].

3. Quantum systems as reservoirs

A novel proposal in this direction is to use quantum systems—particularly quantum computers, whether digital or analogue—as reservoirs. Quantum systems are strong candidates for reservoirs owing to their inherently complex, high-dimensional state spaces and non-trivial dynamics. The underlying idea is straightforward: quantum states live in exponentially large Hilbert spaces, making them natural competitors for generating high-dimensional representations suitable for training a linear readout. Inspired by features emerging from the QML field—as discussed in §3a—a growing body of literature has investigated how to effectively leverage quantum systems as reservoirs, as pioneered by works such as [17,18] and surveyed, e.g. in [19].

In this contribution, we emphasize how it can be especially promising to adopt quantum systems exhibiting strong interactions such as lattices of Rydberg atoms (see §3b). These systems naturally support rich entanglement and strongly correlated behaviour, allowing them to efficiently encode and process temporal correlations in the inputs. Moreover, their tunable interaction strengths and coherence properties enable precise control over the dynamical regime, making it possible to steer the system towards the optimal balance between memory and nonlinearity required for effective RC.

(a) From quantum feature maps to quantum reservoirs

In classical machine learning, *feature mapping* is a well-established concept whereby a certain vector of data $u \in \mathcal{U}$ (the feature)[1] gets mapped via a transformation f from the initial space $\mathcal{U}$ to a different feature one, in order to better exploit some properties of the transformed dataset $f(u) \in \mathcal{F}$. A commonly used example is achieving the linear separability of a dataset for classification purposes, when adopting a support vector machine (e.g. [20]).

A key problem in QML is to load a typically classical input on a quantum device, in order to further process it. The way to achieve this usually invokes a *quantum feature map* acting on a (classical) input u, and targeting a quantum Hilbert space $\mathcal{H}$ such that

$$u \in \mathcal{U} \rightarrow |\psi(u)\rangle\,. \tag{3.1}$$

This mapping is often a quantum circuit parameterized via u, known as the *ansatz*, and acting as a unitary transformation $U_\psi[f(u)]$ of an initial fiducial state, on to the quantum state $|\psi(u)\rangle \in \mathcal{H}$, where we adopted the common *bra–ket* notation. The features u can then be extracted (in part or in full) by measuring the expectation value of an observable as

$$\langle O\rangle = \langle\psi(u)|\, O\, |\psi(u)\rangle\,, \tag{3.2}$$

which implies an additional (nonlinear) transformation of the data. Therefore, one may treat not only the data encoding f but also the associated measurement as the feature map, as the latter is inherited in order to extract any information from a quantum state. Various encodings are possible, affecting the scaling in the width and depth of the quantum circuit required to embed a dataset of given dimensionality, e.g. M copies of size $N \equiv \dim(u)$, but also the suitability of the protocol in near-term quantum devices. A review of the main possible techniques can be found, e.g. in [21], and state-of-the-art methods have produced highly expressive feature maps by appropriately combining classical and quantum operations [20,22].

We illustrate various feature maps (both on the level of unitary encoding and measurement) in figure 1. First, we consider the most basic scheme of a single-qubit system that is transformed with respect to a single-qubit rotation $R_Z(\alpha) = \exp(-i\alpha Z/2)$, where Z is the Pauli-Z matrix, and the initial state $|+\rangle = 1/\sqrt{2}(|0\rangle + |1\rangle)$. Taking Pauli-$X$ as the observable O, the expectation value with respect to the output state $|\psi(a)\rangle = R_Z(a)\,|+\rangle$ is given as $\langle X\rangle = \cos(a)$. Note that the cosine argument

[1]For notation brevity, we use u for both vectors and scalars, and the nature of u is self-explanatory from the context.

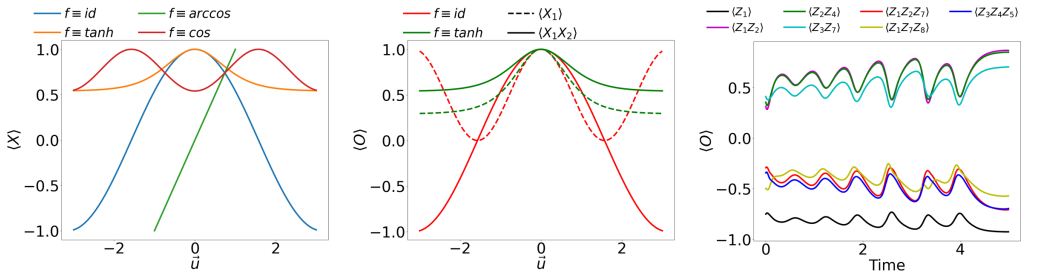

Figure 1. Effect of various feature map encodings. (Left) Single-qubit case with one-dimensional data u on the outcome, expressed as the expectation value of a state $R_Z[f(u)]\,|+\rangle$. (Middle) Two-qubit case, where two types of observables are extracted, as distinguished by dashed versus solid lines. The colours identify different feature map encodings (identity and tanh). (Right) Eight-qubit case encoding a 3D time series, i.e. $\dim(u) = 3$, and values of various colour-encoded observables as a function of time.

can also undergo a feature map transformation as $a \equiv f(u)$, yielding $\langle O\rangle = \cos[f(u)]$. We also illustrate the effect of replacing the identity mapping $f \equiv id$ with various nonlinear functions: $f \equiv \tanh$, $f \equiv cos$ and $f \equiv arccos$.[2] Similarly, in figure 1 (middle), we consider the two-qubit case, where the initial state $|+\rangle\,|+\rangle$ is first rotated with $R_Z(a) \otimes R_Z(a)$ (same Z rotation acting on each qubit) followed by CX gate. In the two-qubit case, we have a more abundant structure of observables that one can measure, and we compare single-qubit expectation value $\langle X_1\rangle$ and two-qubit correlator $\langle X_1X_2\rangle$ with different feature maps encoding the input data. For both single- and two-qubit toy examples, one notices that different encodings transform the data in completely different manners, leaving us with a flexible subroutine to meet machine learning objectives. Finally, for a multi-qubit system and multi-parameter input data (previously $u \equiv u$ was a scalar), the landscape looks more complex. To illustrate this, we plot in figure 1 (right) various observables extracted from encoding input data from the chaotic system Lorenz63, detailed and adopted later in §3c. Here, the dependency is against time, as the Lorenz63 is a three-dimensional (3D) time series, and the encoding phase mixes all components of the input vector together.

When bearing in mind the requirements for a good reservoir structure elicited in §3, a first question concerning equation (3.1) is to what extent we can introduce nonlinearities in a quantum reservoir adopting a feature-mapping scheme? In the classical case, such nonlinearities were naturally present in the dynamical map. However, the evolution of a closed quantum system (an approximation valid for many quantum devices) is linear and described by unitary transformations U. Then, a natural choice is to invoke a nonlinear behaviour within the feature map encoding of the data. That is, $f(u)$ shall be a nonlinear function of u.

Additional challenges and opportunities arise when considering the way we extract classical information from a quantum system—i.e. the measurement. The latter invokes the estimation of *observables*, i.e. special operators acting on the system's Hilbert space. In a quantum reservoir framework, this measurement step preludes the final, linear readout operation in equation (2.3). However, a well-known aspect of quantum measurement is that it collapses the internal state of the system in a non-reversible fashion, producing a classically processable *bitstring* as a result, of the same size as the number of qubits in the device. Each bit therein represents the collapsed value in the computational basis of the corresponding qubit, and average estimates like equation (3.2) require access to many copies of the same state and repeated measurements.

The requirements to build a reservoir are somewhat different from other QML tasks like linear regression or classification, where the reset of the system's memory after each measurement does not pose any fundamental issues. Fading memory requirements conflict with 'forgetting' the process generating the output quantum state—which is, to some extent, the outcome of the measurement process. Therefore, other strategies are required to restore the ability of a quantum

[2]Note that in the case of arccos(·) we operate on a truncated domain.

system to retain memory of past inputs, highlighting a central challenge in QRC design, which we discuss in further detail in §4.

(b) Neutral-atom dynamics

Let us consider a specific example of a quantum system suitable for use as a reservoir: analogue quantum simulators based on neutral atoms, particularly arrays of Rydberg atoms [23]. These systems are currently capable of storing hundreds of physical qubits in vacuum chambers sitting at room temperatures, which can be manipulated by laser addressing of appropriate frequency and engineered pulses. In general, such a system comprising n atoms can be described by a Hamiltonian of the form[3]

$$H(\{\Omega_j\},\{\delta_j\},U,t)=\sum_{j=1}^{n}\left(\Omega_j(t)\,X_j+\delta_j(t)\,n_j\right)+\sum_{i<j}U_{ij}(t)n_in_j, \tag{3.3}$$

where $n_j=(1+Z_j)/2$ is the number operator. The first two parameters Ω (*amplitude*) and δ (*detuning*) are connected with the laser light–atoms interaction, whereas the *coupling* strengths U_{ij} can be tuned independently from each other. In principle, all three parameters are dependent on time t, as pulses can be modulated, and the atom positions can be modified via optical tweezers during the computation, allowing for additional operations [24].

In our tests, we tend to implement less general versions of equation (3.3). The three parameters can often be considered static in time, as some neutral-atom devices have adopted fixed atom arrays throughout the computation (thus making U_{ij} constant), and the modulation of the pulses occurs on time scales fast enough that can be considered constant at least block-wise in terms of overall system evolution. An important simplification is also attained by adopting *global* Ω,δ for the whole array, as this removes the demanding requirement of laser-addressing single atomic sites independently [23]. When including all such simplifications, we are left with a type of transverse field Ising model Hamiltonian [25]:

$$H(\Omega,\delta,U)=\sum_{j=1}^{n}\left(\Omega\,X_j+\delta\,n_j\right)+\sum_{i<j}U_{ij}n_in_j. \tag{3.4}$$

Besides the obvious potential of directly quantum-simulating such models directly on a neutral-atomic system, equation (3.4) can still be expressive enough to implement machine learnable models, as demonstrated, e.g. in [20].

(c) Hybrid quantum–classical reservoir computing with neutral atoms

In the spirit of targeting early experimental implementations on NISQ available quantum devices, we focus now on describing a possible *hybrid* quantum–classical RC (hQCRC) framework, suitable in particular to the neutral-atom platform introduced in §3b. A hybrid approach can serve as a remedy for fading memory [26], as (in its most elementary design) the dynamical map can be split into two parts: classical and quantum. The former handles memory effects, while the latter exploits an exponentially large Hilbert space for the data encoding—with caveats discussed in detail in §3c(i). Both parts process their respective data streams with nonlinear characteristics, thus having the necessary components to form a good RC architecture.

The framework we propose is inspired by [3] (see also [27], which further exploits classical processing techniques), with adjustments to the neutral-atom platform, which will eventually provide digital–analogue capabilities as in [28]. A schematic workflow focusing on sequential datasets is depicted in figure 2. The framework operates in the following three phases:

[3]Hereby, we assume $\hbar=1$ to simplify notation.

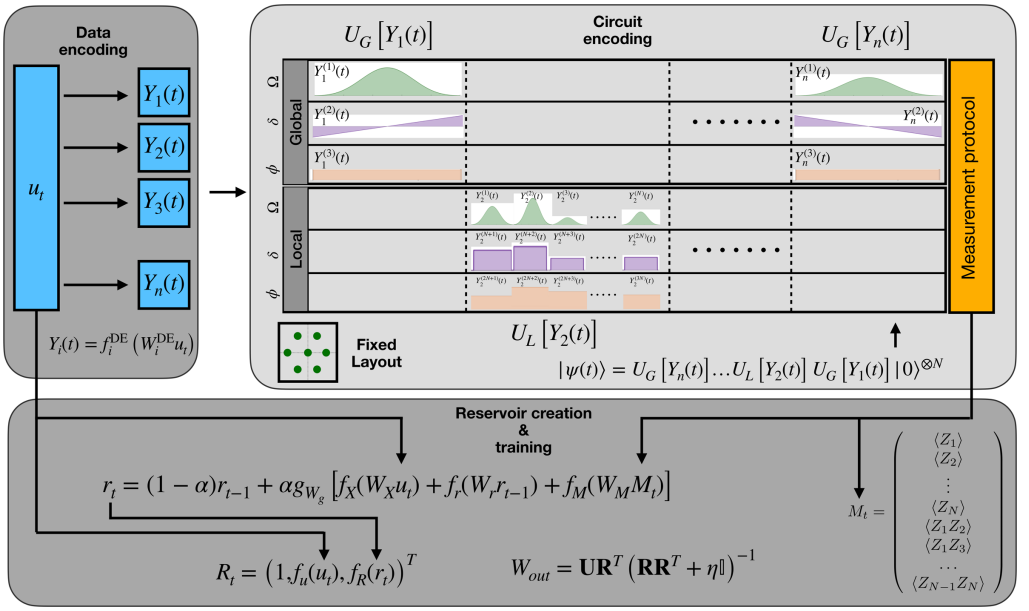

Figure 2. A representation of the three phases of the suggested hybrid QRC protocol. A data encoding manipulates the initial vector u_t into a modified $Y_i(t)$, which is then encoded into the quantum circuit, where i indicates the layer in which the data are injected. The encoding is here exemplified for a neutral-atom architecture compatible with equation (3.3) and processes the input $Y_i^{(j)}(t)$ according to the chosen (sets of) free parameters j such as Ω, δ, ϕ, alternatively operating at a global or local level (details in the main text). The pulse sequences shown for the circuit encoding are generated in `Pulser` [29]. The readout from the quantum circuit M_t is finally post-processed in a reservoir-like fashion compatible with the equation (2.1) and trained until r_t reproduces the target dataset. Phases implemented in classical (quantum) devices are encircled in dark (light) grey boxes.

(i) *Classical data pre-processing*. Responsible for creating suitable input data to feed into the dynamical map as mentioned in §2.

(ii) *Quantum encoding*. A subroutine of the dynamical mapping whose general behaviour we described in equation (2.2).

(iii) *Classical post-processing*. Combining contributions from the classical and quantum dynamical maps, whose output serves training or prediction purposes.

This is evidently a hybrid classical–quantum–classical scheme, relying on various transformations. First (i), the input data are prepared as $\{Y_i(t)\}$ to fit the quantum ansatz, representing a quantum subroutine of the dynamical map, hence it is transformed via combination of linear and nonlinear transformations $f_i^{\mathrm{DE}}: D_{\mathrm{input}} \to D_{\mathrm{layer}}$, where $D_{\mathrm{input/layer}}$ is the dimension of input/layer. Linear transformations are matrix multiplication with fixed—usually random—matrices W_i^{DE} (see also equation (2.3)), whereas the nonlinear ones are activation functions similar to those elicited in figure 1.

The input data are transformed n times with different (random) W_i^{DE} and different f_i^{DE}, where n is the number of parameterized (data-dependent) layers in the quantum circuit (ii). The quantum ansatz is constructed by first fixing the arrangement of atoms,[4] which is responsible for interaction between atoms and thus also the generation of entanglement. Next, a series of laser pulses, algorithmically arranged into layers, is applied to encode the transformed data into a quantum state through time-dependent evolution. Each layer has a number of tunable parameters: Ω, δ, ϕ (described in §3b) as well as the duration (which is fixed per layer in our experiments). The number of

[4]In principle, the arrangement could also depend on transformed data $Y_i(t)$; however, changing the arrangement from iteration to iteration extends wall-clock execution time. Therefore, one needs to take these overheads into account.

terms depends on the type of pulses that are applied, e.g. with local pulses, one has the flexibility to address 1 to $N-1$ different atoms, while global pulses addressing the whole atom register allow for adjustment of only the introduced parameters. Also, in principle, the shape of each pulse can be modelled in a time-dependent fashion; however, constant pulses can still provide sufficient data encoding expressivity. Future generations of neutral-atom platforms are expected to provide even more flexibility; for example, Raman transitions and in-processing layout re-arrangement are areas of active research [23,24].

The quantum state after the full evolution, i.e. when all layers have been applied, has the form

$$|\psi(t)\rangle = U_G\left[Y_n(t)\right] \ldots U_L\left[Y_2(t)\right] U_G\left[Y_1(t)\right] |0\rangle^{\otimes N}, \tag{3.5}$$

where $U_{G/L}[Y_i(t)]$ is a global/local unitary (pulse) parameterized by transformed data $Y_i(t)$. This state is then measured according to a measurement protocol that enables us to create a measurement vector M_t composed of expectation values of selected observables. For example, for Z-type measurements, one may create a measurement vector comprising single-qubit expectation values and two-qubit correlators as

$$M_t = \left(\langle Z_1\rangle, \langle Z_2\rangle \ldots, \langle Z_N\rangle \langle Z_1 Z_2\rangle, \ldots \langle Z_{N-1} Z_N\rangle\right)^T, \tag{3.6}$$

where the ansatz design specifies which qubits and what correlators are used. The goal of a quantum encoding subroutine is to spread classical data in a high-dimensional Hilbert space, while preserving the possibility for measurement protocols to extract such compressed information. In order to expect benefits from the quantum encoding, one needs to identify a protocol that captures the most relevant signatures of the encoded problem. Using exponentially many observables for quantum state tomography is already infeasible for modest system sizes. In general, engineering the choice of observables may require sophisticated techniques, such as adaptive measurement protocols [30] to trade-off between the total number of measurements (different types or bases and the number of shots) and the information extracted. Our implementation focuses on hardware-efficient protocols that match native or easily implementable measurement bases in the chosen neutral-atom architecture.

As a final step (iii), we process available streams of data: input data, measurement vector and previous reservoir states to construct a compact reservoir state r_t:

$$r_t = (1-\alpha) r_{t-1} + \alpha g_{W_g}\left[f_X(W_X u_t) + f_r(W_r r_{t-1}) + f_M(W_M M_t)\right], \tag{3.7}$$

where α is a leak rate, f_X, f_r, f_M are nonlinear functions responsible for transforming the input data u_t, previous reservoir state r_{t-1} and measurement vector M_t, that first undergo linear transformation with fixed (usually random) matrices W_X, W_r and W_M, respectively. The entire second part can additionally be transformed in a similar way, i.e. a combination of linear and nonlinear parts with a function $g_{W_g} \equiv g(W_g \cdot)$. The reservoir state update is responsible for memory effects by mixing previous states with current construction, allowing it to learn time characteristics. The reservoir states are constituents of a reservoir vector R_t:

$$R_t = \left(1, f_u(u_t), f_R(r_t)\right)^T, \tag{3.8}$$

where unity plays the role of a bias, and f_x, f_R are nonlinear activation functions for the input data u_t and the reservoir state r_t,[5] respectively. The reservoir vectors are then aggregated into a matrix $\mathbf{R} = [R_p, R_{p+1}, \ldots, R_d]$, where we discarded the first $p-1$ vectors to avoid reliance on the initial conditions and kept the remaining vectors from the training set ($t = 1, 2, \ldots, d$, keeping $d-p$). The collected reservoir vectors are used to determine an output matrix W_{out} in the ridge regression

[5]In principle, one can combine these nonlinear transformations, with yet another set of linear transformations. Our numerical experiments on small system sizes suggest that it is unnecessary to obtain reasonably good outcomes.

procedure that aims to find the least-square mean between th target data $\mathbf{u} = [u_p, u_{p+1}, \ldots, u_d]$ and the training data $\mathbf{R}$ as

$$W_{out} = \mathbf{U}\mathbf{R}^T \left(\mathbf{R}\mathbf{R}^T + \eta \mathbb{I}\right)^{-1}, \tag{3.9}$$

where η is the regularization coefficient that stabilizes the solution. The output matrix is then used to predict future time steps as

$$y_{t+1} = W_{out} R_t, \quad \text{for} \quad t = d, d+1, \ldots. \tag{3.10}$$

This set-up gives us a flexible and modular approach that can be easily adjusted to hardware restrictions, e.g. the currently cloud-available neutral-atom hardware supports *global* channels only [29,31]. In addition, it provides a good test bed for investigating the transition between classical and QRC. In particular, one can turn off the quantum part by setting $f_M \equiv 0$ or turn off the dependence of the reservoir states on the input data (setting $f_X \equiv 0$) to investigate the importance of quantum or classical contribution. The standard echo state network is recovered for $f_X = f_r = id$, $f_M \equiv 0$ and $g_{W_g} = \tanh$ (e.g. [32]).

As a demonstration of the functionality of the approach, we present in figure 3 a comparison between the predictions from three different incarnations of the hQCRC based on equation (3.7): (i) with all non-trivial terms, i.e. $f_X \equiv f_M \equiv id$; (ii) without the quantum part, i.e. $f_M \equiv 0$; and (iii) with quantum but without classical part, i.e. $f_M \equiv id$ and $f_X \equiv 0$. All cases use $f_r \equiv id$. The (i–iii) approaches are additionally compared with a standard RC (see §2). The results use a chaotic system—Lorenz63—as a benchmark [5], which allows us to assess the performance of the framework in a challenging time-series forecasting task. The hQCRC results use noiseless simulations, based on Schrödinger evolution and numerical precision for determining expectation values. Three different hQCRC approaches use the same ansatz as the one described in figure 2 on an eight-qubit register, with five global constant pulses (layers) of duration 1000 ns, followed by the fixed measurement protocol composed of Z-type measurements to get $\langle Z_i \rangle, \langle Z_i Z_j \rangle$ and $\langle Z_i Z_j Z_k \rangle$ on all possible combinations of i, j, k. We use identity feature maps for global pulses and tanh for f_u and f_R. Further details are reported in an external repository [33]. The purpose of the displayed results is to demonstrate the relevance of the quantum part, as all other parameters are the same, serving as a fair comparison between classical and quantum. Therefore, one may perceive the contribution from the quantum processing as a beneficial addition to the classical system (blue line in figure 3), which is based on exploiting the Hilbert space (circuit) encoding and extracting information from the quantum states via measurements. In addition, one notices that the classical contribution (see green line in figure 3) can be a helpful addition; however, information extracted from the quantum states through the measurement protocol is sufficient to get statistically relevant predictions. Note that we did not optimize over the ansatz hyperparameters and similarly did not use a highly optimized RC implementation. These results are thus intended as a representative co-processing boost attainable via our hQCRC approach on existing analogue neutral-atom platforms. More thorough benchmarks will be required to assess average performances and ultimate benefits of a hybrid protocol (as attained in [3] for digital approaches).

(i) Concentration properties

We have mentioned how one of the principal motivations for QRC is the exponentially increasing dimensionality of the Hilbert space that hosts quantum states. However, as the system size grows, more intricate interactions are required to adequately explore this space, which may lead to dynamics resembling Haar-random evolution. Such behaviour can be detrimental to QML and variational quantum algorithms, as the resulting concentration of measure produces vanishing variances in observables. Consequently, distinguishing between two states differing only in ansatz parameters would demand an exponentially large number of measurements to achieve high-fidelity readout. This phenomenon has been the subject of extensive recent study (see [34] for a review). Here, we present a preliminary analysis of the concentration properties of the hQCRC

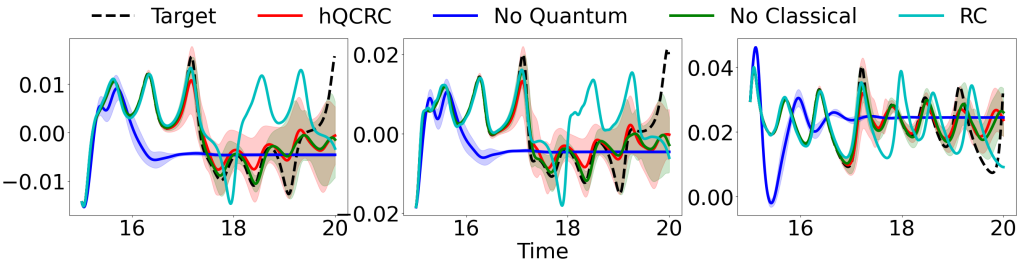

Figure 3. Prediction phase of the normalized Lorenz63 time-series benchmark (black) for x, y and z components (left to right). The red line represents a numerical simulation of the hQCRC framework in equation (3.7) (i.e. $f_M \equiv id$), whereas the blue line synthetically retrieves a classical-phase-only performance by imposing $f_\chi \equiv 0$ in the same equation, while the green line uses $f_\chi \equiv 0$, i.e. the reservoir r_t depends only on measurements extracted from the quantum states M_t. Finally, the teal line depicts the best-performing set-up for classical RC. Implementation details of the circuit are in the main text. The shaded region describes $\pm$ s.d. over 20 random seeds of quantum reservoirs.

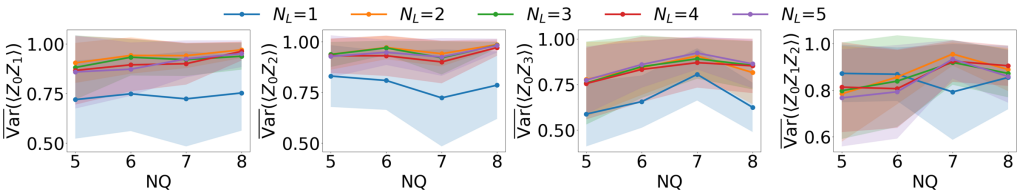

Figure 4. Results for average variance of exemplary two- and three-body correlators over 10 000 random input values to the hQCRC ansatz with varying number of layers (pulses) N_L and number of qubits (NQ). Lines indicate average values, while shaded regions depict spread owing to s.d.

algorithm, deferring a full scaling study to future work. In [35], both numerical and theoretical investigations of reservoir-type quantum extreme learning machines were reported. For multilayer (multi-pulse) implementations, the eigenstructure of the full time-dependent evolution is analytically intractable. Instead, we employ scaling arguments to examine the variance of two- and three-body correlators—commonly used in measurement protocols—as a function of the number of qubits and global pulse layers (each of duration 1000 ns). Figure 4 shows the resulting variances for random inputs uniformly drawn from the interval $[-1, 1]$, where each layer receives independent random values of Ω, δ and phase ϕ. Variances are averaged over 10 000 random initializations. Our results indicate that increasing either the number of qubits or the number of layers yields a relatively stable variance, in contrast to the decay expected from concentration of measure. While the system size and number of layers are limited (the latter by the maximum total evolution time of 6 μs), these results provide initial evidence that programmable atomic arrays may host regimes free from concentration effects. Looking beyond specific QRC architectures, experiments with neutral atoms have demonstrated the presence of quantum many-body scars [36], which may play a key role in preserving learnability, as well as suggested schemes for concentration-free quantum kernels hard to simulate classically [37]. A rigorous analytical and numerical investigation of these mechanisms for RC applications remains an important avenue for future work.

4. The road so far, the road to come

In §3, we elicited some fundamental challenges affecting, in general, the use of quantum systems as performing reservoirs beyond feature mapping—such as the ability to retain fading memory over sequences of inputs, a hallmark of RC. We thus described in §3c how one can defy such criticalities in an exemplary demonstration, focusing on a neutral-atom platform and benefitting from a hybrid architecture to bypass some desirable reservoir properties lacking from what is readily

implementable in the described quantum device. Here, we intend to discuss in perspective some of these design choices and how they might evolve in the near future.

While quantum systems are naturally suited for generating expressive, high-dimensional representations, ensuring that they also possess the temporal dynamics necessary for memory remains non-trivial. We elicited before, the incompatibility between projective quantum measurements and the necessity for the reservoir to retain memory during the processing. A straightforward workaround is to re-initialize and evolve the quantum reservoir from scratch for each new data point, but this quickly becomes intractable as the number of inputs grows. Hybrid approaches like ours, or the one introduced in [26], outsource the memory build-up to classical co-processing. A more advanced approach for mitigating this issue would be to adopt so-called *weak measurements*, which are agnostic to the specific quantum hardware used. Rather than performing a full projective measurement—as illustrated before, which collapses the quantum state entirely—weak measurements couple 'reservoir qubits' to 'ancillary qubits'. Only the latter is targeted to extract classical information, with minimal disturbance to the reservoir quantum state. This strategy implies a careful trade-off between the strength of the measurement, the memory capacity of the reservoir and the computational overhead of resetting the system to successfully perform on benchmark tasks [38,39]. Enhancing feature map techniques, like those suggested above, with mid-circuit measurements can be a promising avenue towards fully fledged QRC, with a particular appeal for neutral-atom devices. In the latter, the availability to use shelving atomic states has led to demonstrations of both ancillary qubits measured mid-circuit [40] as well as entire *zoning* of storage sub-registers for error detection purposes [24]. Bringing these novel techniques to maturity, especially in terms of preserving the fidelity of the involved qubits, will pave the way to increased hybridization of quantum resources in reservoir protocols.

On the other hand, a proper reservoir's behaviour shall prioritize recent inputs and thus needs to erase information over time. Mathematically, it means that the map that describes the reservoir state evolution, r in equation (2.2), needs to be a contraction. As a consequence, the fully unitary dynamics introduced in §3 is not suitable for QRC purposes. Again, in our hQCRC setting, this issue is overcome by post-processing the measured data from the quantum system with a classical fading memory function. However, moving to a fully quantum regime, fading memories could only be obtained if non-negligible losses are present [41]. Within this strategy, an envisioned setting is to evolve an initial quantum state but extract classical information only via measurements on a subsystem. Concretely, for a neutral-atom system compatible with the one in figure 2, one starts with an N-qubit initial state $|0\rangle^{\otimes N}$ and maps, e.g. the transformed input $Y_1(t)$, to a set of Hamiltonian parameters $Y_1^{(1,2,3)}$. Owing to the presence of losses, the evolved state after a fixed time t becomes a more complex function $\mathcal{L}(Y_1^{(1,2,3)}, t)$ generalizing over equation (3.5), where $\mathcal{L}$ is the so-called *Liouvillian* that models the lossy reservoir evolution. The reservoir state r_1 is extracted by weak-measuring a subset of $M < N$ atoms, e.g. repeating the experiments so that finite-shot expectation values of local observables can be computed. For each subsequent input u_t, the procedure can be repeated, but the additional $\mathcal{L}(Y_n^{(1,2,3)}, t)$ acts on the quantum state as evolved until the previous $n-1$ time step and is altered by the weak measurement.

In a more general case, depending on the particular encoding, different physical properties of the quantum system itself can be used for introducing effective losses to the system. For example, if the input is encoded in a reservoir subsystem, scrambling dynamics are favourable for having a fading memory [42]. Also, many open quantum systems are described by a *Markovian* behaviour, which erases information of past inputs exponentially fast over time [43], precluding the possibility of solving tasks where a coexistence of short and long time correlations is needed. This challenge, in a fully quantum case, can only be addressed with the introduction of quantum non-Markovianity [43]. However, hybrid scenarios like the one introduced can leverage classical memory to allow for such longer time correlations.

Another fundamental challenge is the extraction of information from the quantum device, encompassing beyond RC to more generic QML tasks. Data processing is expected to occur in the vast Hilbert space already available to modest-sized quantum devices; however, it is often

lower-dimensional classical data that are required as the usable output for real-world tasks. This throughput mismatch is known as the *readout problem*, and it is closely related to the aforementioned concentration phenomena. Specific techniques are being explored to extract such crux of information [44], often at the heart of the performance of quantum algorithms. In our specific case, this readout challenge emerges naturally, e.g. when adopting ancillary subsystems and scrambling dynamics, because the number of measurements needed to properly compute the reservoir response can exponentially scale with the reservoir dimension [45,46], and only the presence of appropriate symmetries can solve the issue [46]. A measurement overhead problem can also be present in the neutral-atom setting considered here, as the system size scales and (linearly) more expectation values need to be computed in equation (3.6). Possible strategies for dealing with this issue will be the topic of further research.

On a positive note, QRC does not require expensive circuit gradient estimate and optimization [11], which affects instead other variational QML techniques. Therefore, it could elude trainability issues connected with the emergence of complex or flat optimization landscapes [34], while harnessing the expressivity accessible to complex, high-dimensional quantum states to potentially capture patterns in specific datasets better than classical tools [26].

Finally, we can safely assume that noise limitations in the quantum device of choice for QRC implementations will play a crucial role in dictating architectures. At a stage in QC where extremely different technologies coexist, to the point that appropriate cross-platform benchmarking is still an actively developed and debated research topic [1,47], it is expected that different quantum architectural choices—with their strengths and weaknesses in terms of gate fidelity, scalability in circuit width and depth, repetition rate and access to native operations—will influence the design of the corresponding QRC protocols. Indeed, each of these aspects will affect noise differently on the processing and readout phases. In this contribution, we focused on a neutral-atom implementation, but even restricting ourselves to a single platform and its related noise model, we can envisage how a careful design and optimization of the ansatz can bring very significant performance benefits. The long-term complexity of this optimization problem, as circuits scale beyond classical simulability and ease of result interpretation, will eventually invoke the aid of automated strategies, in a similar manner to that which has occurred in the field of variational QML with quantum architecture search schemes [48].

5. Conclusions

We have presented a vision for QRC implemented in the hQCRC framework. The purpose of this approach is to mitigate all shortcomings of a quantum system for being a potent reservoir, while taking advantage of the larger latent training space accessible to such systems.

In this regard, we showed that (quantum) feature maps represent an important tool for RC, as they enable processing in such high-dimensional latent space. We mitigated the effect of destructive measurements by moving the memory part of the reservoir from the quantum subroutine to a classical one. We also combined the encoding of input data into quantum states through analogue evolution (data-to-circuit), with nonlinearities associated with the measurement protocol. In addition, our hQCRC framework uses feature maps on the level of classical reservoir states, making it more flexible to accommodate complex tasks. The balance between purely classical and hybrid processing in our approach can be tuned, and the quantum part can be easily turned completely off to study the effect of the latter on task performances.

Crucially, we highlighted how our hQCRC can fit current and future hardware. We demonstrated how the mapping described above can invoke arrays of interacting qubits, where reconfigurable Rydberg neutral-atom arrays represent a prominent example of such a quantum architecture. The modularity of the quantum circuit construction enables us to scale the complexity of each component in the framework—type and number of operations addressing the qubits, the size of the register and the number of extracted observables (which translates into reservoir sizes). Therefore, one can adjust the hQCRC approach to meet hardware restrictions. Notably,

the presented framework exhibits an overarching structural form that can be suitably adjusted to meet hardware requirements and still be efficiently implementable. This means that one can accommodate hQCRC to gate-based, ion traps, neutral atoms or even annealers with marginal changes in the ansatz.

We demonstrated in a proof-of-concept experiment how the hQCRC implemented on a simulated neutral-atom platform can predict future evolution of chaotic time-series system, and the quantum subroutine can be effectively exploited to augment the performance of a classical counterpart. Within the scope of this experiment, we did not observe any immediate trainability issues of the chosen architecture, which might, though, appear in much larger instances and prompt appropriate taming countermeasures. Our results show how the efficient combination of quantum and classical resources can be beneficial for information processing in typical RC tasks, such as time-series prediction. Our work opens a route to new quantum–classical RC approaches and raises the question of useful resources in the hybrid domain.

Data accessibility. The chaotic dataset Lorenz63 is in the public domain. The code to run and reproduce the numerical experiments in the text will be made available under the Zenodo repository [49].

Declaration of AI use. We have not used AI-assisted technologies in creating this article.

Authors' contributions. C.G.: formal analysis, investigation, methodology, validation, writing—original draft, writing—review and editing; F.W.: conceptualization, data curation, formal analysis, investigation, methodology, software, validation, visualization, writing—review and editing; E.J.P.: data curation, investigation, resources, software; A.S.: methodology, validation, writing—review and editing; H.S.: conceptualization, project administration; O.K.: investigation, supervision, validation, writing—review and editing; D.V.: conceptualization, investigation, supervision, validation, writing—review and editing; A.A.G.: conceptualization, investigation, supervision, writing—original draft, writing—review and editing.

All authors gave final approval for publication and agreed to be held accountable for the work performed therein.

Conflict of interest declaration. Pasqal SaS is a for-profit company. Technology and tools referenced in this work could be patent-pending or copyrighted by Pasqal, as detailed in the references provided.

Funding. C.G., O.K. and A.A.G. acknowledge the support from the European Union's Horizon Europe research and innovation programme under Grant Agreement 101070347-MANNGA. O.K. acknowledges support from the QCi3 Hub (grant number EP/Z53318X/1). A.S. acknowledges the Spanish State Research Agency, through the María de Maeztu project CEX2021-001164-M funded by the MCIU/AEI/10.13039/501100011033, through the COQUSY project PID2022-140506NB-C21 and -C22 funded by MCIU/AEI/10.13039/501100011033, MINECO through the QUANTUM SPAIN project, and EU through the RTRP—NextGenerationEU within the framework of the Digital Spain 2025 Agenda. A.S. also acknowledges the CSIC Interdisciplinary Thematic Platform (PTI+) on Quantum Technologies in Spain (QTEP+) and the support of a fellowship from the 'la Caixa' Foundation (ID100010434—LCF/BQ/DI23/11990081). C.G., O.K. and A.A.G. acknowledge the support from the European Union's Horizon Europe research and innovation programme under Grant Agreement 101070347-MANNGA.

Acknowledgements. Views and opinions expressed are those of the authors only and do not necessarily reflect those of the European Union or the European Health and Digital Executive Agency (HADEA). Neither the European Union nor the granting authority can be held responsible for them.

References

1. Proctor T, Young K, Baczewski AD, Blume-Kohout R. 2025 Benchmarking quantum computers. *Nat. Rev. Phys.* **7**, 105–118. (doi:10.1038/s42254-024-00796-z)
2. Callison A, Chancellor N. 2022 Hybrid quantum-classical algorithms in the noisy intermediate-scale quantum era and beyond. *Phys. Rev. A* **106**, 010101. (doi:10.1103/PhysRevA.106.010101)
3. Wudarsk F, OConnor D, Geaney S, Asanjan AA, Wilson M, Strbac E, Lott PA, Venturelli D. 2023 Hybrid quantum-classical reservoir computing for simulating chaotic systems. *arXiv*. (doi:10.48550/arXiv.2311.14105)
4. Huang HY, Choi S, McClean JR, Preskill J. 2025 The vast world of quantum advantage. *arXiv*. (doi:10.48550/arXiv.2508.05720)

5. Platt JA, Penny SG, Smith TA, Chen TC, Abarbanel HDI. 2022 A systematic exploration of reservoir computing for forecasting complex spatiotemporal dynamics. *Neural Netw.* **153**, 530–552. (doi:10.1016/j.neunet.2022.06.025)
6. Yan M, Huang C, Bienstman P, Tino P, Lin W, Sun J. 2024 Emerging opportunities and challenges for the future of reservoir computing. *Nat. Commun.* **15**, 2056. (doi:10.1038/s41467-024-45187-1)
7. Marković D, Mizrahi A, Querlioz D, Grollier J. 2020 Physics for neuromorphic computing. *Nat. Rev. Phys.* **2**, 499–510. (doi:10.1038/s42254-020-0208-2)
8. Gauthier DJ, Bollt E, Griffith A, Barbosa WAS. 2021 Next generation reservoir computing. *Nat. Commun.* **12**, 5564. (doi:10.1038/s41467-021-25801-2)
9. Cucchi M, Abreu S, Ciccone G, Brunner D, Kleemann H. 2022 Hands-on reservoir computing: a tutorial for practical implementation. *Neuromorph. Comput. Eng.* **2**, 032002. (doi:10.1088/2634-4386/ac7db7)
10. Jaeger H, Haas H. 2004 Harnessing nonlinearity: predicting chaotic systems and saving energy in wireless communication. *Science* **304**, 78–80. (doi:10.1126/science.1091277)
11. Jaeger H. 2001 *The 'echo state' approach to analysing and training recurrent neural networks-with an erratum note*. p. 148. Bonn, Germany: German NRC for information technology - technical report.
12. Gallicchio C *et al*. 2019 Chasing the echo state property. In *ESANN 2019-Proc., 27th European Symposium on Artificial Neural Networks, Computational Intelligence and Machine Learning*. Bruges, Belgium: ESANN.
13. Zhong Y, Tang J, Li X, Gao B, Qian H, Wu H. 2021 Dynamic memristor-based reservoir computing for high-efficiency temporal signal processing. *Nat. Commun.* **12**, 408. (doi:10.1038/s41467-020-20692-1)
14. Finocchio G, Di Ventra M, Camsari KY, Everschor-Sitte K, Khalili Amiri P, Zeng Z. 2021 The promise of spintronics for unconventional computing. *J. Magn. Magn. Mater.* **521**, 167506. (doi:10.1016/j.jmmm.2020.167506)
15. Van der Sande G, Brunner D, Soriano MC. 2017 Advances in photonic reservoir computing. *Nanophotonics* **6**, 561–576. (doi:10.1515/nanoph-2016-0132)
16. Penkovsky B, Larger L, Brunner D. 2018 Efficient design of hardware-enabled reservoir computing in FPGAs. *J. Appl. Phys* **124**, 162101. (doi:10.1063/1.5039826)
17. Fujii K, Nakajima K. 2021 Quantum reservoir computing: a reservoir approach toward quantum machine learning on near-term quantum devices. In *Reservoir computing: theory, physical implementations, and applications* (eds K Nakajima, I Fischer), pp. 423–450. Singapore: Springer Singapore. (doi:10.1007/978-981-13-1687-6_18)
18. Marković D, Grollier J. 2020 Quantum neuromorphic computing. *Appl. Phys. Lett* **117**, 150501. (doi:10.1063/5.0020014)
19. Mujal P, Martínez-Peña R, Nokkala J, García-Beni J, Giorgi GL, Soriano MC, Zambrini R. 2021 Opportunities in quantum reservoir computing and extreme learning machines. *Adv. Quantum Tech* **4**, 2100027. (doi:10.1002/qute.202100027)
20. Albrecht B *et al*. 2023 Quantum feature maps for graph machine learning on a neutral atom quantum processor. *Phys. Rev. A* **107**, 042615. (doi:10.1103/PhysRevA.107.042615)
21. Schuld M, Killoran N. 2019 Quantum machine learning in feature Hilbert spaces. *Phys. Rev. Lett.* **122**, 040504. (doi:10.1103/PhysRevLett.122.040504)
22. Kyriienko O, Paine AE, Elfving VE. 2021 Solving nonlinear differential equations with differentiable quantum circuits. *Phys. Rev. A* **103**, 052416. (doi:10.1103/PhysRevA.103.052416)
23. Henriet L, Beguin L, Signoles A, Lahaye T, Browaeys A, Reymond GO, Jurczak C. 2020 Quantum computing with neutral atoms. *Quantum* **4**, 327. (doi:10.22331/q-2020-09-21-327)
24. Bluvstein D *et al*. 2024 Logical quantum processor based on reconfigurable atom arrays. *Nature* **626**, 58–65. (doi:10.1038/s41586-023-06927-3)
25. Labuhn H, Barredo D, Ravets S, de Léséleuc S, Macrì T, Lahaye T, Browaeys A. 2015 Realizing quantum ising models in tunable two-dimensional arrays of single rydberg atoms. *arXiv.*
26. Kornjača M *et al*. 2024 Large-scale quantum reservoir learning with an analog quantum computer. *arXiv.* (doi:10.48550/arXiv.2407.02553)
27. Settino J *et al*. 2025 Memory-augmented hybrid quantum reservoir computing. *Phys. Rev. Appl.* **24**, 024019. (doi:10.1103/wzwv-7rk2)

28. Parra-Rodriguez A, Lougovski P, Lamata L, Solano E, Sanz M. 2020 Digital-analog quantum computation. *Phys. Rev. A* **101**, 022305. (doi:10.1103/PhysRevA.101.022305)
29. Silvério H, Grijalva S, Dalyac C, Leclerc L, Karalekas PJ, Shammah N, Beji M, Henry LP, Henriet L. 2022 Pulser: an open-source package for the design of pulse sequences in programmable neutral-atom arrays. *Quantum* **6**, 629. (doi:10.22331/q-2022-01-24-629)
30. Chen SYC, Tseng HH, Lin HY, Yoo S. 2025 Learning to measure quantum neural networks. In *2025 IEEE Int. Conf. on Acoustics, Speech, and Signal Processing Workshops (ICASSPW)*, pp. 1–5. Hyderabad, India: IEEE. (doi:10.1109/ICASSPW65056.2025.11011001)
31. Pasqal. 2025 Pasqal SaS Cloud Services. See https://docs.pasqal.com/cloud/.
32. Lukoševičius M. 2012 A practical guide to applying echo state networks. In *Neural networks: tricks of the trade* (eds G Montavon, GB Orr, KR Müller), pp. 659–686, 2nd edn. Berlin, Germany: Springer. (doi:10.1007/978-3-642-35289-8_36)
33. Wudarski F, Evan John P, Antonio Andrea G. 2025 From quantum feature maps to quantum reservoir computing: an applicative perspective. Zenodo. (doi:10.5281/zenodo.17153722)
34. Larocca M *et al*. 2025 Barren plateaus in variational quantum computing. *Nat. Rev. Phys*. **7**, 174–189. (doi:10.1038/s42254-025-00813-9)
35. Xiong W, Facelli G, Sahebi M, Agnel O, Chotibut T, Thanasilp S, Holmes Z. 2025 On fundamental aspects of quantum extreme learning machines. *Quantum Mach. Intell.* **7**, 20. (doi:10.1007/s42484-025-00239-7)
36. Bernien H *et al*. 2017 Probing many-body dynamics on a 51-atom quantum simulator. *Nature* **551**, 579–584. (doi:10.1038/nature24622)
37. Sarkar A, Schnee M, Radgohar R, Fadaie M, Drouin-Touchette V, Kourtis S. 2025 Concentration-free quantum kernel learning in the Rydberg Blockade. *arXiv*. (doi:10.48550/arXiv.2508.10819)
38. Mujal P, Martínez-Peña R, Giorgi GL, Soriano MC, Zambrini R. 2023 Time-series quantum reservoir computing with weak and projective measurements. *Npj Quantum Inf* **9**, 16. (doi:10.1038/s41534-023-00682-z)
39. Franceschetto G, Płodzień M, Lewenstein M, Acín A, Mujal P. 2024 Harnessing quantum back-action for time-series processing. *arXiv*. (doi:10.48550/arXiv.2411.03979)
40. Graham TM *et al*. 2023 Midcircuit measurements on a single-species neutral alkali atom quantum processor. *Phys. Rev. X* **13**, 041051. (doi:10.1103/PhysRevX.13.041051)
41. Sannia A, Martínez-Peña R, Soriano MC, Giorgi GL, Zambrini R. 2024 Dissipation as a resource for quantum reservoir computing. *Quantum* **8**, 1291. (doi:10.22331/q-2024-03-20-1291)
42. Martínez-Peña R, Giorgi GL, Nokkala J, Soriano MC, Zambrini R. 2021 Dynamical phase transitions in quantum reservoir computing. *Phys. Rev. Lett*. **127**, 100502. (doi:10.1103/PhysRevLett.127.100502)
43. Sannia A, Rodríguez RR, Giorgi GL, Zambrini R. 2025 Non-markovianity and memory enhancement in quantum reservoir computing. *arXiv*. (doi:10.48550/arXiv.2505.02491)
44. Williams CA, Scali S, Gentile AA, Berger D, Kyriienko O. 2024 Addressing the readout problem in quantum differential equation algorithms with quantum scientific machine learning. *arXiv*. (doi:10.48550/arXiv.2411.14259)
45. Xiong W, Holmes Z, Angrisani A, Suzuki Y, Chotibut T, Thanasilp S. 2025 Role of scrambling and noise in temporal information processing with quantum systems. *arXiv*. (doi:10.48550/arXiv.2505.10080)
46. Sannia A, Giorgi GL, Zambrini R. 2025 Exponential concentration and symmetries in quantum reservoir computing. *arXiv*. (doi:10.48550/arXiv.2505.10062)
47. Lorenz JM *et al*. 2025 Systematic benchmarking of quantum computers: status and recommendations. *arXiv*. (doi:10.48550/arXiv.2503.04905)
48. Martyniuk D, Jung J, Paschke A. 2024 Quantum architecture search: a survey. In *2024 IEEE Int. Conf. on Quantum Computing and Engineering (QCE)*, vol. 1, pp. 1695–1706, Montréal, Canada: IEEE. (doi:10.1109/QCE60285.2024.00198)
49. Wudarski F, Philip EJ, Gentile AA. 2025 From quantum feature maps to quantum reservoir computing: an applicative perspective. Zenodo. (doi:10.5281/zenodo.17153722)