PROCEEDINGS OF SCIENCE

Quantum Reservoir Learning on Pasqal neutral-atom Quantum Computer

Simone Angelozzi,[a,b,*] **Simone Gasperini**[a,c] **and Daniele Bonacorsi**[a,c]

[a] *Department of Physics and Astronomy "Augusto Righi", University of Bologna, 40127 Bologna, Italy*
[b] *School of Chemistry, University of Southampton, Highfield, Southampton SO17 1BJ, United Kingdom*
[c] *National Institute for Nuclear Physics (INFN), Sezione di Bologna, 40127 Bologna, Italy*

E-mail: s.angelozzi@soton.ac.uk, simone.gasperini4@unibo.it

Neutral-atom quantum devices offer a natural and powerful platform for Quantum Reservoir Computing (QRC), a non-variational approach to quantum-enhanced machine learning that is attracting growing interest in the Noisy Intermediate-Scale Quantum (NISQ) era.
In this work, we implement a QRC-based learning model tailored to Pasqal neutral-atom hardware, where classical data are encoded into time-dependent detuning waveforms and processed through programmable many-body dynamics. Feature embeddings are extracted from single- and two-qubit Pauli-Z observables and used within a simple linear readout model. We systematically investigate both global and local detuning encoding strategies on a supervised image classification task, benchmarking performance against a purely classical linear baseline. Our results show that processing the data through the quantum reservoir yields a consistent ~4% improvement in accuracy, highlighting the system effectiveness as a nonlinear feature map. Both encoding strategies exhibit comparable performance, suggesting that QRC can leverage different control schemes while remaining compatible with near-term and emerging hardware capabilities.
These findings provide concrete evidence that QRC can enhance classical machine learning pipelines by injecting non-classical correlations without the need for variational optimization, thereby avoiding issues such as barren plateaus and training instabilities. More broadly, this work establishes neutral-atom platforms as a promising and scalable setting for QRC, opening the way toward practical quantum-enhanced learning on near-term devices.

International Symposium on Grids and Clouds (ISGC2026)
15-20 March, 2026, Academia Sinica, Taipei, Taiwan

*Speaker

© Copyright owned by the author(s) under the terms of the Creative Commons Attribution-NonCommercial-NoDerivatives 4.0 International License (CC BY-NC-ND 4.0) All rights for text and data mining, AI training, and similar technologies for commercial purposes, are reserved. ISSN 1824-8039 . Published by SISSA Medialab.

PoS(ISGC2026)017

1. Introduction

The rapid progress in quantum technologies, together with the advent of Noisy Intermediate-Scale Quantum (NISQ) computers, has inaugurated a new computational paradigm characterized by systems of tens to hundreds of qubits with limited connectivity [1]. Although this scale already challenges classical simulation, the presence of errors and noise continues to limit the practical applicability of such devices.

At the same time, advances in artificial intelligence have spurred growing interest in the intersection of quantum computing and machine learning. On one hand, quantum devices have the potential to accelerate specific machine learning subroutines [2]; on the other hand, classical machine learning techniques can be leveraged to enhance quantum computation [3]. This mutual interplay naturally motivates the development of hybrid quantum-classical approaches that aim to exploit the strengths of both paradigms.

In the NISQ era, the dominant framework for Quantum Machine Learning (QML) is represented by Variational Quantum Algorithms (VQAs), which rely on parameterized quantum circuits trained via classical optimization routines, typically based on gradient descent [4]. Despite promising proof-of-concept results, VQAs face several fundamental challenges: the emergence of barren plateaus due to the variational structure and the exponentially large Hilbert space, as well as the detrimental effects of hardware noise and local minima [5].

A compelling alternative that has recently attracted attention is Quantum Reservoir Computing (QRC), which mitigates these limitations by eliminating the need for variational parameter optimization [6–8]. Instead of relying on an iterative quantum-classical feedback loop, QRC employs a quantum system with fixed dynamics as a computational reservoir. Input data are encoded into the reservoir, and the resulting observables generated during the system evolution in the Hilbert space are processed by a classical readout model. This approach intrinsically avoids barren plateaus and offers an efficient framework for implementing machine learning tasks on NISQ devices.

Neutral-atom quantum platforms are particularly well suited for QRC implementations due to their programmable Rydberg interactions and flexible addressing capabilities [9, 10]. In this work, we implement a QRC-based classifier for a Pasqal neutral-atom device, encoding classical features through tunable detuning parameters. We systematically investigate two encoding strategies, global and local detuning, and demonstrate that QRC-based models may effectively enrich the input data representation through nonlinear feature expansion. These results highlight QRC as a promising approach for quantum-enhanced machine learning on short-term NISQ hardware.

2. Background

Reservoir Computing is a machine learning framework that circumvents the costly training of internal network parameters by employing a three-layer architecture consisting of an input layer, a fixed dynamical system (the reservoir), and an output layer. This framework exploits the nonlinear transformations induced by the reservoir while restricting training to the output layer, which is typically implemented via simple linear regression. Originally introduced in seminal works on Echo State Networks and Liquid State Machines [11–13], Reservoir Computing was proposed as an efficient alternative to the training of recurrent neural networks. More recently, it has been

PoS(ISGC2026)017

validated as a powerful and versatile machine learning approach due to its simplicity and broad range of practical applications [14, 15].

The extension of Reservoir Computing to quantum systems naturally addresses the training challenges associated with variational quantum algorithms. The most direct formulation follows the hybrid quantum-classical paradigm, in which classical data are encoded into a quantum substrate that processes the information for downstream classical learning tasks [16]. Notably, recent large-scale implementations on neutral-atom systems have demonstrated QRC with up to a hundred qubits using programmable Rydberg interactions [6]. These studies successfully applied QRC to supervised learning tasks such as time-series forecasting and multi-class image classification, establishing neutral-atom platforms as a competitive setting for practical QML without the training overhead inherent to variational approaches.

More recent works further highlight the versatility of QRC across diverse application domains. QRC implemented on neutral-atom hardware has been used to predict the biological activity of drug molecules, yielding more robust and interpretable predictions with reduced variance compared to classical baselines [17]. Beyond molecular property prediction, QRC has also been applied to forecasting chaotic systems and to predicting excited-state properties of molecules from ground-state wavefunctions [18–20].

3. Methods

This section outlines the task considered in this work and the overall workflow adopted for QRC implementation. The primary objective is to evaluate the performance of QRC on a supervised machine learning task. We used the MNIST-like dataset available through scikit-learn[1], containing grayscale images of handwritten digits, each represented by an 8×8 grid of pixels. This dataset provides a widely adopted low-resolution benchmark for image classification. The task is formulated as a standard supervised classification problem, where each image is labeled with a digit from 0 to 9, and the ten classes are nearly balanced. Directly encoding all the 64 input features into the QRC model is impractical given current hardware constraints and the computational cost associated with simulating high-dimensional quantum dynamics. To address this limitation, we apply Principal Component Analysis (PCA) to reduce the dimensionality of the data. Each image is first flattened into a 64-dimensional vector, and PCA is fitted on the flattened images of the full dataset, retaining the $N = 5$ leading principal components. This choice matches the number of atoms in the register, enabling a one-to-one mapping between features and atoms in the local encoding scheme while keeping the emulation cost manageable. The PCA-reduced data are split into training (70%) and test (30%) sets with a stratified split that preserves the class proportions. Finally, each principal component is rescaled to the interval $[0, 1]$ through a min-max normalization fitted on the training set only; the same transformation is applied to the test set, clipping possible out-of-range values to the interval boundaries. The resulting five-dimensional feature vectors $\mathbf{x} = (x_1, \ldots, x_5)$ constitute the input of both the QRC models and the classical baseline. The QRC algorithm is implemented using Pulser[2], a Python framework for pulse-level programming on Pasqal neutral-atom devices [21]. This framework enables the definition of quantum registers through the selection of atoms from

[1] https://scikit-learn.org/stable/
[2] https://pulser.readthedocs.io/en/stable/

PoS(ISGC2026)017

the device geometry, the construction of time-dependent pulses to encode data into Hamiltonian parameters, and the classical emulation of quantum dynamics via QuTiP[3] [22].

In the following, we describe the key components of the QRC implementation in detail, including register design, quantum evolution, and observables extraction. All the results reported in this work are obtained through classical emulation of the quantum dynamics, and no experiment has been executed on a physical quantum processor yet. The emulation setting, including how the characteristics of the real *Fresnel* device by Pasqal are taken into account, is detailed in Sec. 3.4.

3.1 Register of atoms

The quantum reservoir is realized by preparing a register composed of a linear chain of 5 equally-spaced atoms, with an inter-atomic distance $a = 5\,\mu m$, as shown in Fig. 1. The spatial arrangement of the atoms is designed to ensure controlled interactions while remaining compatible with the constraints of the neutral-atom platform; in the global encoding scheme, the atoms are placed in traps of the pre-calibrated triangular layout of the device, as required for execution on real hardware. A key parameter governing the system dynamics is the Rydberg blockade radius, which is set to $R_b = 7\,\mu m$. Since R_b exceeds the nearest-neighbor distance but not the next-nearest-neighbor one, blockade effects involve only nearest-neighbor pairs, although the weaker interactions between all atom pairs are included in the emulated dynamics. This determines the effective range of inter-atomic interactions, enabling the generation of entanglement across the register, an essential resource for the reservoir computational capability. Additionally, the blockade radius constrains the choice of the Rabi frequency used in the pulse sequences, thereby influencing the encoding of classical input data into the quantum system according to the selected strategy.

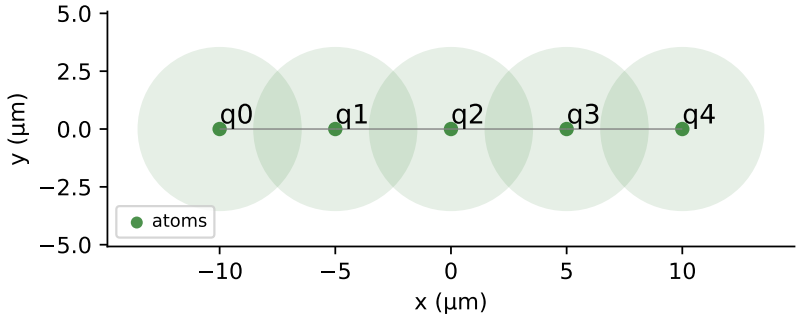

Figure 1: Register of atoms for the QRC algorithm. Green dots represents qubits, light green circles denote the interaction range between them, defined by the blockade radius.

3.2 Sequence of pulses

Following the QRC framework introduced in [6], we exploit the tunability of the Rydberg Hamiltonian to embed classical data into the controllable parameters of the reservoir. The time-dependent Hamiltonian governing the Pasqal neutral-atom platform is

$$\hat{H}(t) = \frac{\hbar\Omega(t)}{2}\sum_j \hat{\sigma}_j^x - \sum_j \hbar[\delta^{\text{glob}}(t) + \delta_j^{\text{loc}}(t)]\hat{n}_j + \sum_{\langle j,k\rangle} \frac{C}{|x_j - x_k|^6}\hat{n}_j\hat{n}_k, \tag{1}$$

[3] https://qutip.org/

PoS(ISGC2026)017

where $\Omega(t)$ is the global Rabi frequency, $\delta^{\text{glob}}(t)$ is the global detuning, and $\delta_j^{\text{loc}}(t)$ is the site-dependent local detuning. The last term accounts for Van der Waals interactions between neighbor atoms $\langle j,k\rangle$ separated by a distance $|x_j - x_k|$. These parameters enable different encoding strategies,

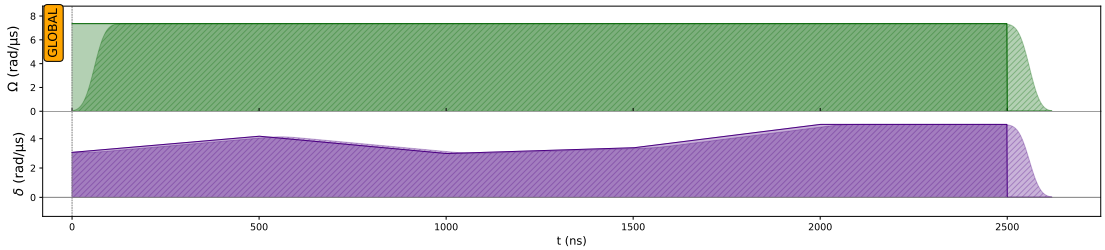

Figure 2: Pulse sequence used in the global detuning encoding scheme. The green line represents the constant Rabi frequency Ω while the purple curve is the time-dependent global detuning δ^{glob}, with classical input features encoded in its tunable amplitude.

which directly affect the expressivity of the reservoir by shaping how data are embedded into the nonlinear dynamics of the system. In this work, we consider:

- **Global detuning encoding**: input features are encoded sequentially into the time-dependent amplitude of the global detuning $\delta^{\text{glob}}(t_i)$, where t_i labels discrete time steps (see the example in Fig. 2). This approach leverages the intrinsic time evolution of the system to generate feature correlations. Notably, the number of features is not limited by the size of the reservoir, but only by the total available evolution time. Moreover, this scheme is compatible with the current Pasqal platform.

- **Local detuning encoding**: input features are encoded in parallel into the constant amplitudes of site-dependent detunings $\delta_j^{\text{loc}}(t)$, enabling a direct mapping between features and atoms. This can enhance expressivity by distributing information spatially across the register, potentially leading to richer many-body dynamics. However, it requires individual atom addressing, which is not currently available on Pasqal devices, and is therefore investigated only via ideal emulation on an unconstrained virtual device (see Sec. 3.4).

Another possible encoding strategy, not explored in this work, is *position encoding*, where classical information is embedded in the spatial configuration of the atomic register. While potentially increasing expressivity by altering connectivity, this approach requires fine-grained control over atom positions and is beyond current hardware capabilities.

3.3 Quantum observables

Information is extracted from the quantum reservoir by measuring a set of observables that define the classical feature space for the learning task. Within the QRC framework, these observables provide a nonlinear mapping of the classical input data, reflecting how it is processed by the underlying quantum dynamics. We consider both single-qubit and two-qubit observables evaluated at discrete time steps. Single-qubit measurements correspond to the expectation values of the operators $\langle Z_j\rangle$, which provide local information at each atom site. To capture interaction-induced

PoS(ISGC2026)017

effects, we also include two-qubit observables $\langle Z_j Z_k \rangle$, encoding pairwise correlations arising from the Rydberg interactions between neighbor atoms.

For each input sample, the system is evolved under the corresponding pulse sequence, and the observables are measured at multiple time steps during the evolution. At a given time t_i, the embedding is constructed by concatenating all single-qubit expectation values with the set of two-qubit correlations between interacting pairs. The final feature vector is then obtained by stacking these embeddings over the full evolution time window $[0, T]$, thereby incorporating both spatial and temporal information from the reservoir. The resulting embedding has dimension $(N_{\text{atoms}} + N_{\text{pairs}}) \cdot N_T$, where N_{atoms} is the number of atoms, N_{pairs} is the number of interacting pairs, and N_T is the number of time steps. In our setting, two-qubit correlators are evaluated for the $N_{\text{pairs}} = 4$ nearest-neighbor pairs lying within the blockade radius, so that with $N_{\text{atoms}} = 5$ and $N_T = 5$ each input sample is mapped into a 45-dimensional embedding. This high-dimensional feature vector is then used as input to the classical readout layer for the supervised learning task.

3.4 Emulation setting

All the quantum dynamics are computed by exact numerical integration of the Schrödinger equation for the five-atom register with the QuTiP backend of Pulser, and the observables are evaluated as exact expectation values on the emulated quantum state. In this context, it is useful to distinguish three levels of emulation: (i) an *ideal emulation*, where the Hamiltonian dynamics is computed without any restriction on the control parameters; (ii) a *device-constrained ideal emulation*, where the pulse sequence is additionally validated against the specifications of a real device; (iii) a *noisy emulation*, where hardware imperfections are explicitly modeled.

The global detuning experiments belong to the second category. The corresponding pulse sequences are built on `AnalogDevice`, the Pulser device specification representative of Pasqal analog QPUs such as *Fresnel*, and are subject to the same validation required for submission to the QPU. This enforces, among others, the use of a pre-calibrated register layout (a triangular lattice of traps with 5 μm spacing), a Rydberg channel with global addressing only, bounds on the Rabi frequency ($\Omega \leq 12.6$ rad/μs) and on the detuning ($|\delta| \leq 125.7$ rad/μs), and the finite time resolution of the control waveforms. The local detuning experiments, instead, belong to the first category, since individual addressing requires the Pulser `MockDevice`, a virtual device without hardware constraints.

In neither case hardware noise is included: state preparation and measurement errors, laser amplitude and phase fluctuations, Doppler dephasing, atom losses, the finite modulation bandwidth of the control fields, and the shot noise arising from a finite number of measurement repetitions are not modeled. The reported results should therefore be regarded as noiseless benchmarks of the proposed models, rather than as a prediction of their performance on current hardware.

3.5 Trainable classical layer

After constructing embedding feature vectors from the quantum reservoir, we evaluate their discriminative power on the supervised classification task. In the QRC framework, the classical readout layer is the only trained component, in line with the Reservoir Computing paradigm where the complexity of the model is shifted to the fixed dynamical system, and learning is restricted to a

PoS(ISGC2026)017

simple output mapping. This stage employs a simple classifier to assess how effectively the quantum embeddings capture task-relevant features, and to benchmark performance against a purely classical baseline. In both cases, we use multinomial logistic regression, a linear classifier that extends binary logistic regression to the multiclass setting. This choice is particularly suitable in the QRC context, as it allows us to directly evaluate the quality of the quantum feature map without introducing additional nonlinearities beyond those generated by the quantum reservoir.

To optimize classifier performance, we perform a grid search with 5-fold cross-validation over the regularization hyper-parameters. The parameters configuration yielding the highest cross-validation accuracy is selected and subsequently evaluated on the test set. Model performance is evaluated using standard classification metrics for balanced supervised tasks, including overall accuracy and confusion matrices.

4. Results

We first test the QRC-based classifier using global detuning encoding, where input features are mapped into the temporal profile of a global detuning pulse $\delta^{\text{glob}}(t)$ applied uniformly to the whole system of atoms. Experiments are performed by noiseless classical emulation, enforcing

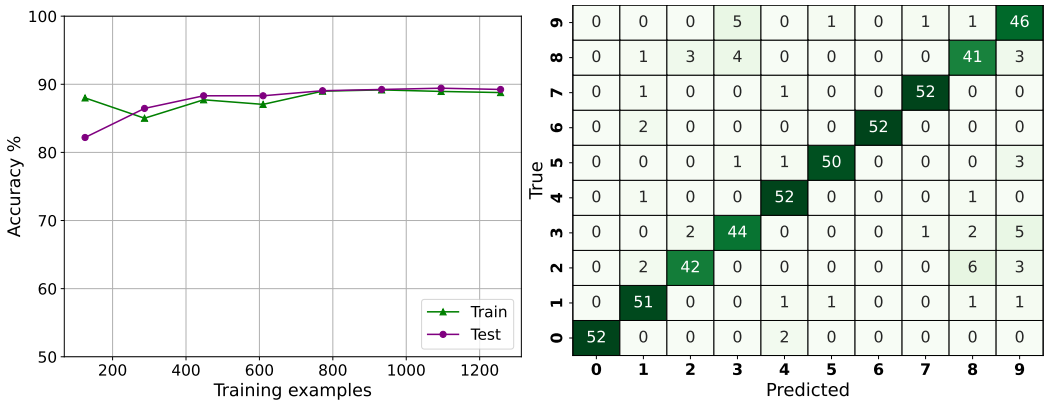

Figure 3: QRC model performance with global detuning encoding. (Left) Learning curve evaluated on a fixed test set, reaching a peak accuracy slightly below the 90%. (Right) Confusion matrix on the fixed test images for the model with the highest test accuracy; rows correspond to true labels and columns to predicted labels.

the physical constraints of the Pulser device specification representative of the *Fresnel* device (see Sec. 3.4) to ensure compatibility with Pasqal real hardware. The Rabi frequency is kept constant at $\Omega \simeq 7.36$ rad/μs, the value corresponding to the blockade radius $R_b = 7\ \mu m$, and the five PCA-reduced features are encoded as detuning amplitudes $\delta^{\text{glob}}(t_i) = \Omega x_i$. The detuning waveform consists of $N_T = 5$ segments of duration $t = 500\ ns$, yielding a total evolution time $T = 2500\ ns$: in each of the first four segments the detuning is linearly ramped from Ωx_i to Ωx_{i+1}, while in the last one it is held constant at Ωx_5. The observables are recorded at the end of each segment.

Results are shown in Fig. 3. The model reaches a test accuracy slightly below the 90%, with a decreasing gap between train and test performance as the number of examples used for training grows, indicating good generalization. The confusion matrix, shown for the model with the highest

PoS(ISGC2026)017

test accuracy along the learning curve, exhibits a strong diagonal dominance. Its rows correspond to the true labels and its columns to the predicted labels. Therefore the matrix is not expected to be symmetric, and the ratio between a diagonal entry and the corresponding row total sum gives the per-class recall. The largest off-diagonal entries involve the pairs 2/8 and 3/9, while digits often regarded as visually similar, such as 1 and 7, are almost never confused.

We also investigate the alternative encoding strategy where each input feature is mapped into the amplitude of the local detuning applied to each atom individually. Unlike global encoding, local addressing requires control over individual atoms, a capability not yet available on Pasqal current hardware. For this reason, local encoding is studied exclusively via ideal classical emulation. For

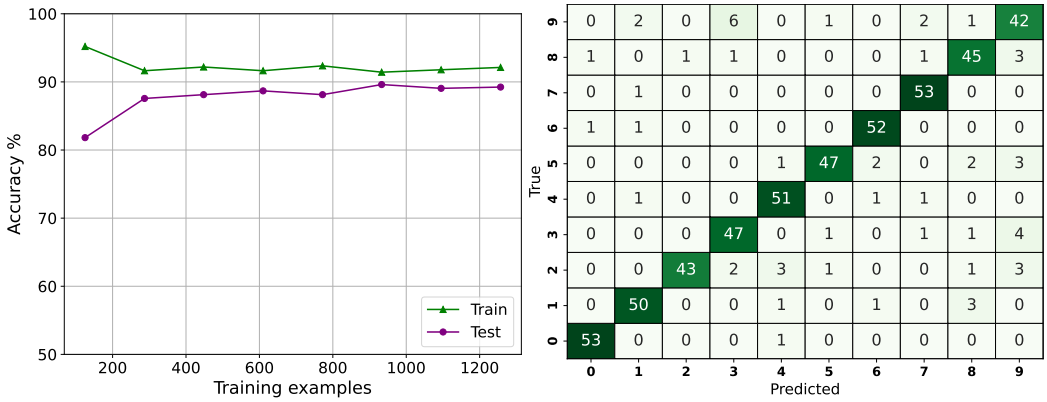

Figure 4: QRC model performance with local detuning encoding. (Left) Learning curve evaluated on a fixed test set, reaching a peak accuracy slightly below the 90%. (Right) Confusion matrix on the fixed test images for the model with the highest test accuracy; rows correspond to true labels and columns to predicted labels.

each input sample, we construct pulse sequences where distinct and site-dependent local detunings δ_j^{loc} are applied in parallel to each atom. Since individual addressing is not supported by the `AnalogDevice` specification, these sequences are built on the virtual `MockDevice` (see Sec. 3.4), using the same register geometry, blockade radius, and set of observables as in the global scheme. Each atom is driven by a constant pulse of duration $T = 500\ ns$, with Rabi frequency $\Omega \simeq 46.1$ rad/μs, corresponding to $R_b = 7\ \mu m$ for the Rydberg level of this device, and local detuning $\delta_j^{\text{loc}} = \delta_{\max} x_j$, with $\delta_{\max} \simeq 48.7$ rad/μs. The observables are recorded at $N_T = 5$ equally spaced instants, every 100 ns. Although the absolute time scales differ from the global scheme, the two settings are dynamically comparable: in both cases the ratio between the nearest-neighbor interaction and the Rabi frequency is $(R_b/a)^6 \simeq 7.5$, the detuning amplitudes span a range of the order of Ω, and the total pulse area ΩT is of the same order.

Results are shown in Fig. 4. The learning curve shows both train and test accuracy as a function of training set size. The model reaches a peak test accuracy slightly below the 90%, with train accuracy consistently higher across all data regimes, indicating stable generalization. Notably, the overall performance of global and local encoding are comparable, despite the different control mechanisms. The confusion matrix, again shown for the model with the highest test accuracy and following the same conventions as in Fig. 3, exhibits a strong diagonal dominance. However, the pattern of residual errors only partially matches the global encoding case. The 3/9 pair remains

PoS(ISGC2026)017

the most frequent source of confusion, whereas the 2/8 pair, prominent with global encoding, is almost absent, and the remaining errors are spread over several pairs. Also in this case, 1 and 7 are essentially never confused.

These observations indicate that misclassifications are not simply determined by the visual similarity of the digits. Indeed, the classifier does not operate on digit shapes as perceived by a human observer, but on reservoir embeddings of the five leading principal components of 8×8 intensity images, which mainly capture the coarse spatial distribution of the ink. In this representation, 1 and 7 are plausibly well separated, since a thin, mostly vertical central stroke and a pronounced upper horizontal bar produce clearly different intensity patterns. Conversely, digits such as 3 and 9, or 2 and 8, share intensity over large common regions of the image and differ mainly in finer details, such as open or closed loops, which are largely washed out by the low resolution and by the truncation to five principal components. The fact that the 3/9 confusion appears with both encodings suggests that it is largely inherited from the reduced input representation, whereas the 2/8 confusion, observed only with global encoding, reflects the specific way in which each reservoir reshapes the feature space.

It is worth stressing that the comparison between the two encoding strategies is not a comparison between an ideal and a noisy emulation. Both schemes are emulated without any source of hardware noise, and they differ only in that the global scheme is additionally validated against the constraints of a real device, while the local scheme relies on the unconstrained virtual device required by individual addressing (see Sec. 3.4). Moreover, neither encoding systematically outperforms the other. Along the learning curve, the test accuracies of the two schemes differ by at most 1%, and both reach the same accuracy of 89.2% when the full training set is used.

The comparable performance is consistent with the fact that, in both cases, the achievable accuracy is mainly bounded by the information retained in the five PCA components and by the linear readout, rather than by the control constraints: removing the latter does not by itself increase the information available to the classifier. In addition, neither the reservoir parameters nor the readout hyper-parameters were optimized separately for each encoding, so that the two schemes are compared under matched, rather than individually optimized, conditions.

To rigorously assess the impact of the quantum reservoir, we compare the performance of the QRC embeddings against a purely classical linear baseline. The objective of this comparison is to isolate the effect of the quantum processing step by evaluating how the nonlinear evolution of the reservoir enhances the capabilities of a standard linear classifier. To this end, the classical baseline consists of the same logistic regression algorithm trained directly on the PCA-reduced input features. Conversely, the QRC models process the input data through the quantum reservoir before applying the linear classifier, effectively employing the physical system as a nonlinear quantum feature extractor.

Figure 5 illustrates the test accuracy of the three approaches. The comparison reveals that processing the data through the quantum reservoir before the linear classification step yields a consistent improvement across all evaluated data regimes. The comparison is carried out in terms of overall test accuracy, which is evaluated for all three models along the same learning curve. Specifically, the classical baseline saturates at a test accuracy of 84.8% with the full training set, whereas both QRC models reach the 89.2%. The gain is smallest for the smallest training set and lies between approximately 3% and 4.5% for larger training sets, with both global and

PoS(ISGC2026)017

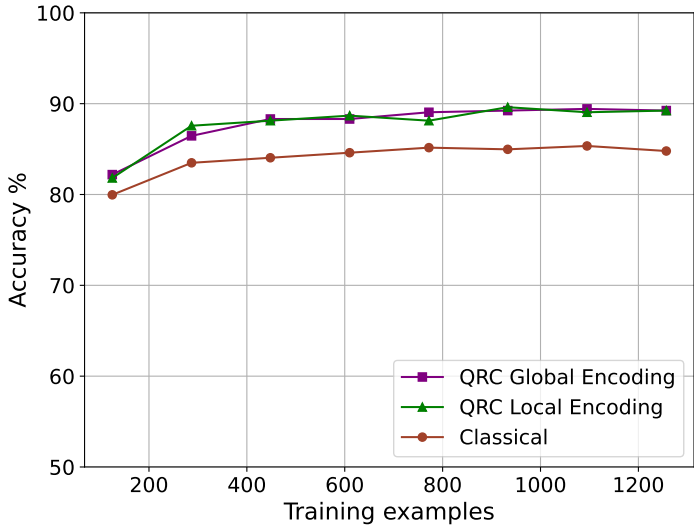

Figure 5: Test accuracy comparison between QRC with global detuning encoding, QRC with local detuning encoding, and the classical linear baseline (logistic regression) as a function of training set size.

local encoding. Since the classifier, its hyper-parameters, and the train/test split are identical in all cases, this improvement can be attributed to the nonlinear feature expansion provided by the quantum reservoir. Moreover, the training accuracy of the baseline remains close to its test accuracy, indicating that the linear model is limited by the separability of the PCA-reduced features rather than by the amount of training data, which is precisely the limitation addressed by the reservoir. Given that tunable time-dependent global detuning is already possible on current Pasqal neutral-atom hardware, these findings highlight its potential viability on current quantum devices.

5. Conclusion

In this work, we presented a Quantum Reservoir Computing approach for supervised image classification, specifically designed considering the constraints and capabilities of the Pasqal neutral-atom platform. By employing the quantum system as a fixed, nonlinear feature extractor, the QRC framework restricts the training process entirely to a classical linear readout layer, thereby circumventing the optimization challenges typical of Variational Quantum Algorithms.

Through noiseless classical emulation, we systematically investigated global and local detuning encoding strategies, incorporating the physical constraints of the Fresnel device in the former case. Our results demonstrate that processing the input data through the quantum reservoir prior to a standard linear classifier yields a consistent 4% improvement on the accuracy, compared to a purely classical linear baseline. This confirms the capability of the quantum dynamics to enrich the input data representation through nonlinear feature expansion. Furthermore, the performance of the global detuning strategy was shown to be highly comparable to the local addressing scheme. This establishes the global detuning approach as an effective and viable encoding strategy for near-term hardware architectures. While these simulated results highlight the potential of neutral-atom QRC, transitioning to physical devices execution remains a critical next step.

PoS(ISGC2026)017

Future work will focus on addressing current hardware implementation challenges, such as the necessity for intermediate partial measurements during reservoir evolution. In parallel, a systematic investigation of the impact of realistic hardware noise and finite shot statistics on the reservoir expressivity will be essential to fully assess the robustness of the QRC paradigm on near-term devices.

In this respect, the noise models available in Pulser, covering effects such as state preparation and measurement errors, laser noise, and dephasing, provide a natural next step before execution on physical hardware.

Although the implementation presented here is tailored to Pasqal neutral-atom hardware, the QRC scheme itself is not specific to this platform. Its essential ingredients are a quantum system with fixed and sufficiently rich dynamics, a controllable parameter through which classical data can be injected, and a set of measurable observables [16]. On superconducting quantum computers, for instance, the detuning encoding could be mapped onto data-dependent single-qubit Z rotations or qubit-frequency modulations, the Rydberg interactions could be replaced by fixed entangling layers approximating an Ising-type evolution, and the same single- and two-qubit Z correlators could be extracted. Neutral-atom platforms remain particularly convenient in this context, since the reservoir dynamics is analog and its interaction graph can be reconfigured through the atom positions, whereas on gate-based devices the reservoir would be shaped by the native connectivity and affected by the errors accumulated in deeper circuits. The present study is restricted to the ten-class handwritten digit recognition task, which serves as a controlled benchmark for the proposed QRC pipeline. A natural extension would be to consider datasets including both letters and digits, focusing on visually ambiguous character pairs such as B and 8 or Z and 2, in order to probe the ability of the reservoir to resolve fine-grained differences between classes and to further assess the broader applicability of the approach.

Overall, our findings highlight QRC as a promising framework for integrating complex quantum nonlinearities into hybrid machine learning pipelines. In this context, an additional compelling direction would be the application of QRC-based approaches to time-series tasks, for which this framework appears particularly well suited due to its natural ability to process temporal data.

References

[1] J. Preskill, *Quantum Computing in the NISQ era and beyond*, *Quantum* **2** (2018) 79.

[2] M. Klusch, J. Lässig, D. Müssig, A. Macaluso and F. K. Wilhelm, *Quantum Artificial Intelligence: A Brief Survey*, *KI - Künstliche Intelligenz* **38** (2024) 257.

[3] G. Carleo, I. Cirac, K. Cranmer, L. Daudet, M. Schuld, N. Tishby et al., *Machine learning and the physical sciences*, *Reviews of Modern Physics* **91** (2019) 045002.

[4] M. Cerezo, A. Arrasmith, R. Babbush, S. C. Benjamin, S. Endo, K. Fujii et al., *Variational quantum algorithms*, *Nature Reviews Physics* **3** (2021) 625.

[5] M. Larocca, S. Thanasilp, S. Wang, K. Sharma, J. Biamonte, P. J. Coles et al., *Barren plateaus in variational quantum computing*, *Nature Reviews Physics* **7** (2025) 174.

PoS(ISGC2026)017

[6] M. Kornjača, H.-Y. Hu, C. Zhao, J. Wurtz, P. Weinberg, M. Hamdan et al., *Large-scale quantum reservoir learning with an analog quantum computer*, July, 2024. 10.48550/arXiv.2407.02553.

[7] Q. Li, C. Mukhopadhyay, L. Minati and A. Bayat, *Quantum reservoir computing for predicting and characterizing chaotic maps*, *Phys. Rev. Res.* **8** (2026) 013304.

[8] O. Ahmed, F. Tennie and L. Magri, *Robust quantum reservoir computers for forecasting chaotic dynamics: generalized synchronization and stability*, *Proceedings of the Royal Society A: Mathematical, Physical and Engineering Sciences* **481** (2025) 20250550.

[9] L. Henriet, L. Beguin, A. Signoles, T. Lahaye, A. Browaeys, G.-O. Reymond et al., *Quantum computing with neutral atoms*, *Quantum* **4** (2020) 327.

[10] M. Saffman, T. G. Walker and K. Molmer, *Quantum information with Rydberg atoms*, *Reviews of Modern Physics* **82** (2010) 2313.

[11] M. Lukoševičius and H. Jaeger, *Reservoir computing approaches to recurrent neural network training*, *Computer Science Review* **3** (2009) 127.

[12] H. Jaeger, *The "echo state" approach to analysing and training recurrent neural networks*, Tech. Rep. GMD Report 148, German National Research Center for Information Technology, Bonn, Germany, 2001.

[13] W. Maass, T. Natschläger and H. Markram, *Real-time computing without stable states: a new framework for neural computation based on perturbations*, *Neural Computation* **14** (2002) 2531.

[14] K. Nakajima, H. Hauser, T. Li and R. Pfeifer, *Information processing via physical soft body*, *Scientific Reports* **5** (2015) 10487.

[15] A. Deihimi and H. Showkati, *Application of echo state networks in short-term electric load forecasting*, *Energy* **39** (2012) 327.

[16] P. Mujal, R. Martínez-Peña, J. Nokkala, J. García-Beni, G. L. Giorgi, M. C. Soriano et al., *Opportunities in Quantum Reservoir Computing and Extreme Learning Machines*, *Advanced Quantum Technologies* **4** (2021) 2100027.

[17] D. Beaulieu, M. Kornjača, Z. Krunic, M. Stivaktakis, J. Chen, T. Ehmer et al., *Robust quantum reservoir learning for molecular property prediction*, *Journal of Chemical Information and Modeling* **65** (2025) 8475.

[18] O. Ahmed, F. Tennie and L. Magri, *Robust quantum reservoir computers for forecasting chaotic dynamics: generalized synchronization and stability*, *Proc. R. Soc. A* **481** (2025) 20250550.

[19] O. Ahmed, F. Tennie and L. Magri, *Prediction of chaotic dynamics and extreme events: A recurrence-free quantum reservoir computing approach*, Oct., 2024. 10.48550/arXiv.2405.03390.

PoS(ISGC2026)017

[20] H. Kawai and Y. O. Nakagawa, *Predicting excited states from ground state wavefunction by supervised quantum machine learning*, *Machine Learning: Science and Technology* **1** (2020) 045027.

[21] H. Silvério, S. Grijalva, A. Cornillot, L. Henriet, A. Quelle, L. AJDNIK et al., *pasqal-io/pulser: v1.7.2*, Mar., 2026. 10.5281/zenodo.18889313.

[22] N. Lambert, E. Giguère, P. Menczel, B. Li, P. Hopf, G. Suárez et al., *QuTiP 5: The quantum toolbox in Python*, 2026. https://doi.org/10.1016/j.physrep.2025.10.001.

PoS(ISGC2026)017