

Measurements as a design principle for quantum reservoirs: enhancing and enabling online temporal learning

This extended abstract is based on two recent works: *Harnessing Quantum Backaction for Time-Series Processing*, published in *Physical Review X* [1], and *Measurement-enabled online quantum processing with amplitude encoding*, available as an arXiv preprint [2].

Classical and quantum reservoir computing. Reservoir computing provides a training-efficient paradigm for temporal information processing. In classical machine learning, recurrent architectures such as Long Short-Term Memory networks have become standard tools for sequence modelling, but their performance relies on the optimization of large sets of internal parameters. Reservoir computing follows a different route: the input sequence drives a fixed dynamical system, whose complex response generates a feature space for the task. Only the final readout is trained, typically by a simple linear regression. In this way, the expressive power required for temporal processing is delegated to the intrinsic dynamics of the reservoir, while the learning stage remains simple [3, 4]. This idea is particularly appealing in quantum machine learning, where training large parametrized quantum models is challenging both algorithmically and experimentally [5, 6]. Quantum reservoir computing (QRC) transfers the reservoir-computing philosophy to quantum systems by using a quantum dynamical system as the fixed processor [7–9]. The input stream is encoded into the reservoir, information is spread through its natural or engineered evolution, and measured observables provide the classical feature vector used for inference. Since only the output layer is optimized, QRC offers a near-term route to quantum temporal learning that avoids training the full quantum dynamics while retaining the possibility of exploiting complex many-body evolution available in experimental platforms.

What changes in the quantum setting. This translation from classical to quantum reservoirs is not automatic. A useful reservoir computer must generate a rich input–output map while retaining information about recent inputs and progressively forgetting irrelevant past history. In classical reservoirs, nonlinearity provides a rich input–output map, while contractive or leaky dynamics make the influence of past inputs progressively fade. By contrast, the evolution of a closed quantum system is linear and unitary: it preserves distinguishability and does not by itself provide the contraction mechanism usually associated with fading memory. As a result, QRC requires an input–update mechanism that introduces an effective nonlinearity in the dependence on the classical data and, at the same time, supplies the memory structure needed for temporal processing.

A standard way to address this difficulty is the amplitude-encoding update introduced in the original formulation of QRC [7], which has since become a reference model for quantum temporal processing [10–13]. At each time step, a subset of the reservoir is traced out and replaced by a new quantum state parametrized by the incoming classical input. This trace-and-replace operation plays a double role: it introduces a nonlinear dependence on the input data and provides a contraction mechanism that supports fading memory. At the same time, it separates the input injection from the subsequent reservoir evolution, making it possible to study the internal quantum dynamics as the processing step of the algorithm. Although this update is natural at the level of the theoretical QRC map, realizing it as an online physical protocol is nontrivial. At each input step, the encoding subsystem must be traced out and reinitialized in an input-dependent state, while the remaining qubits continue to store the reservoir memory. This operation has to be repeated without restarting the dynamics and, ideally, while still allowing access to observables from the full reservoir. Therefore, a gap remains between the theoretical formulation of amplitude-encoded QRC and its implementation on current quantum platforms.

A second limitation arises from the quantum nature of the readout. Temporal processing requires the reservoir to retain information about previous inputs while new inputs are being processed. At the same time, the output layer is trained from observables extracted from the quantum system at each time step. These two requirements are in tension: projective measurements provide information about the reservoir, but they also disturb the state in which temporal memory is stored. Consequently, information extraction is not a passive operation in online QRC, but a dynamical ingredient that can affect the subsequent evolution of the reservoir. Restarting protocols avoid accumulated measurement disturbance by reprocessing the input sequence before each readout, while rewinding protocols reduce this overhead by replaying only a finite window of recent inputs [13, 14]. These strategies protect the reservoir from the direct accumulation of projective-measurement backaction, but they compromise the online character of the computation by requiring repeated preparation or input replay. Online protocols based on indirect measurements take a different route: they monitor the reservoir during a single forward evolution, leading to a runtime that scales linearly with the number of time steps [1, 13]. However, they introduce an information–disturbance tradeoff, since weaker measurements preserve more memory but extract less information, whereas stronger measurements provide more informative readout at the cost of greater perturbation.

Our contribution: measurements as a design principle. We address these obstacles by taking measurements as a design principle for QRC. First, we show analytically that mid-circuit measurement and reset provide an experimentally friendly way to realize the partial trace required by amplitude encoding: the input subsystem is refreshed at each time step, while the memory subsystem continues to carry the reservoir state forward [2]. Second, for the




FIG. 1: Measurement-enabled online quantum reservoir computing. (a) Dynamic-circuit implementation of amplitude-encoded QRC. The scheme shows the k -th block of the protocol with amplitude encoding on the input register R and indirect measurement on the memory register M . (b) Measurement back-action as a resource. The color maps show performance-ratio analyses as a function of the reservoir parameter h and the indirect-measurement strength g for temporal-processing benchmarks (Forecasting on Santa Fe, above, and Short term memory below). The insets summarize the finite-resource analysis. (c) Hardware demonstration of the online amplitude-encoded protocol. Forecasting and short-term-memory capacities are shown as a function of the prediction distance η , comparing ideal spin-model simulations, circuit simulations, and superconducting-hardware data.

readout problem, we do not treat indirect measurements only as weak probes used to limit disturbance. We treat their strength as a tunable parameter of the reservoir and show that an optimal measurement regime can enhance temporal processing [1]. Finally, we show how this tunable indirect monitoring can be implemented with dynamic circuits, using ancilla-mediated measurements to access memory-qubit observables without projectively measuring the memory [2]. These ingredients converge into an online QRC protocol that combines amplitude encoding, tunable indirect readout, linear runtime in the number of time steps, and a proof-of-principle demonstration on superconducting quantum hardware.

Measurement-enabled amplitude encoding. We first solve the implementation gap of amplitude-encoded QRC. In the theoretical model, the input update is defined as a trace-and-replace operation: before the k -th reservoir evolution, the encoding subsystem is traced out and reinitialized with the next input, while the remaining subsystem keeps the reservoir memory. Let R denote the input-reset register, M the memory register, and $\tilde{\rho}_{k+1}$ the joint state after the reservoir evolution at time step k . The amplitude-encoding update assumes the map

$$\tilde{\rho}_{k+1} = \mathcal{E}^{(R,M)} \left[\rho^{(R)}(s_{k+1}) \otimes \text{Tr}_R(\tilde{\rho}_k) \right], \quad (1)$$

where $\rho^{(R)}(s_{k+1})$ is the state encoding the next input s_{k+1} , and $\mathcal{E}^{(R,M)}$ is the reservoir evolution coupling R and M . The question is how to realize the non-unitary operation inside the brackets on hardware. We show that it is exactly implemented by a dynamic-circuit step: measure R , reset it, and prepare the next input state. If the measurement-reset operation is described by $\Omega_{i_k} = |0\rangle\langle i_k|_R \otimes I_M$, then

$$\mathcal{J}_{k+1}^{(R)} \left[\sum_{i_k} \Omega_{i_k} \tilde{\rho}_k \Omega_{i_k}^\dagger \right] = \rho^{(R)}(s_{k+1}) \otimes \text{Tr}_R(\tilde{\rho}_k). \quad (2)$$

where $\mathcal{J}_{k+1}^{(R)}$ prepares the reset register in the input state $\rho^{(R)}(s_{k+1})$. Equations (1) and (2) show that the formal trace-and-replace update of amplitude encoding can be realized as a hardware-friendly dynamic-circuit primitive. The input register is measured, reset, re-prepared, and reused at each time step, while the memory register carries temporal information forward.

The resulting protocol is a single-pass implementation of amplitude-encoded QRC. At each time step, the input register R is prepared with the current data point s_k , coupled to the memory register M , measured, reset, and reused for the next input. The memory register is not reinitialized, and therefore carries the temporal state of the reservoir across the sequence. This avoids the repeated reconstruction of the dynamics required by restarting or rewinding strategies: the circuit advances once through the data stream, with a linear runtime. In this sense, mid-circuit measurement and resets provide the basic dynamic-circuit structure for an online, amplitude-encoded quantum reservoir.

This measurement-reset block is illustrated in Fig. 1(a), which shows the input register being prepared, coupled to the memory register, measured, reset, and reused while the memory register propagates the reservoir state across time.

Back-action as a resource. The same online structure also makes clear where the readout problem remains. Observables on the input register can be obtained directly, since this register is reset at every time step. By contrast, the memory register stores the temporal state of the reservoir, and projectively measuring it would interrupt the subsequent dynamics. Restricting the readout to the reset register would leave part of the reservoir inaccessible, while directly measuring the memory register would compromise the online evolution. We therefore use ancillary qubits to monitor the memory register indirectly: an ancilla is coupled to a memory qubit, measured, and reset, while the memory qubit remains in the recurrent dynamics. This provides access to memory-register observables without turning the readout into a destructive projective measurement of the reservoir state.

Crucially, the role of indirect measurements is not only to make readout less destructive. In online QRC, the measurement strength controls an information–disturbance tradeoff: weak measurements preserve more of the reservoir state but provide less informative observables, while stronger measurements extract more information at the cost of larger backaction. We show that this tradeoff should be optimized rather than minimized. By scanning the measurement strength g together with the reservoir parameters, we identify regimes in which the online protocol with indirect measurements outperforms protocols following the unperturbed reservoir dynamics [1]. This is quantified through performance-ratio analyses. The ratio P_R compares the temporal-processing capacity of the online protocol with indirect measurements to that of an unperturbed restarting protocol at the same reservoir parameter; therefore, $P_R > 1$ means that the measured online dynamics performs better than the corresponding unperturbed one. As summarized in Fig. 1(b), the optimal performance occurs at finite measurement strength, with regions where $P_R > 1$. These results show that measurement backaction is not merely tolerated, but can become a computational resource for temporal processing. The advantage also persists beyond the ideal infinite-sampling limit. Under finite measurement resources, the maximal performance ratio remains above the $P_R = 1$ threshold and approaches the ideal value as the number of measurements (Inner panels of Fig 1 (b)). We further compare the optimized indirect-measurement protocol with a feedback-driven strategy in which previous expectation values are stored classically and re-injected through a unitary feedback layer [15]. In the short-term memory task, the optimized indirect-measurement protocol maintains higher performance at longer delays, indicating that preserving quantum memory through controlled back-action can outperform strategies based only on classical feedback.

A hardware-friendly online protocol. The same online circuit block that realizes amplitude encoding through measurement and reset of the input register can be extended with ancilla-mediated measurements of the memory register (Fig. 1 (a)). In this way, the protocol combines the trace-and-replace input update, memory preservation, and tunable indirect readout in a single repeated circuit. The input register is refreshed at every time step, the memory register propagates the temporal state, and the ancillary probes provide access to memory observables without directly measuring the memory qubits. This gives a hardware-compatible implementation of online amplitude-encoded QRC in which measurement is not an external readout operation, but part of the recurrent dynamics itself. We validate this measurement-designed protocol in a 9-qubits proof-of-principle experiment on superconducting quantum hardware. The implementation uses dynamic-circuit primitives, namely mid-circuit measurement and reset, to realize the repeated online block on the `ibm_basquecountry` Heron r2 processor. Since the spin-reservoir evolution is not native to the device, the entangling dynamics are implemented through a hardware-tailored Trotterized circuit. We benchmark the resulting reservoir on two standard temporal tasks of more than one thousand data points, Santa Fe time-series forecasting and short-term memory [4, 16, 17], and compare hardware data with both ideal spin-model simulations and circuit simulations. As shown in Fig. 1(c), the hardware results reproduce the ideal capacity trends and almost exactly match those of the noisy simulated protocol. This demonstrates that amplitude encoding, memory preservation, and indirect online monitoring can be combined in a dynamic-circuit implementation of QRC. The protocol relies only on mid-circuit measurement and reset operations, which are widely available in present quantum devices [18–21].

General takeaway. Taken together, these results support a measurement-designed view of online QRC. Measurement is not only the final step used to extract classical data from a quantum reservoir. It can implement the input-update rule, preserve the memory subsystem during online processing, provide tunable access to internal observables, and even enhance temporal performance when its back-action is optimized. The relevant object is therefore the full recurrent channel realized by the experiment: input preparation, reservoir evolution, measurement, reset, and post-measurement state update. This perspective connects the theoretical requirements of amplitude-encoded QRC with experimentally available dynamic-circuit primitives, and provides a route toward online quantum reservoirs in which measurement is used as an architectural resource rather than treated only as a limitation.

-
- [1] G. Franceschetto, M. Płodzień, M. Lewenstein, A. Acín, and P. Mújal, Harnessing quantum backaction for time-series processing, *Phys. Rev. X* **16**, 021002 (2026).
 - [2] G. Franceschetto, P. Mújal, and R. Martínez-Peña, Measurement-enabled online quantum processing with amplitude

- encoding (2026), arXiv:2606.18991 [quant-ph].
- [3] H. Jaeger and H. Haas, Harnessing nonlinearity: Predicting chaotic systems and saving energy in wireless communication, *Science* **304**, 78 (2004).
 - [4] H. Jaeger, *Short Term Memory in Echo State Networks*, Tech. Rep. GMD Report 152 (GMD Forschungszentrum Informationstechnik GmbH, 2002).
 - [5] J. Bowles, S. Ahmed, and M. Schuld, Better than classical? the subtle art of benchmarking quantum machine learning models (2024), arXiv:2403.07059 [quant-ph].
 - [6] M. Larocca, S. Thanasilp, S. Wang, K. Sharma, J. Biamonte, P. J. Coles, L. Cincio, J. R. McClean, Z. Holmes, and M. Cerezo, Barren plateaus in variational quantum computing, *Nature Reviews Physics* **7**, 174–189 (2025).
 - [7] K. Fujii and K. Nakajima, Harnessing disordered-ensemble quantum dynamics for machine learning, *Phys. Rev. Appl.* **8**, 024030 (2017).
 - [8] S. Ghosh, K. Nakajima, T. Krisnanda, K. Fujii, and T. C. H. Liew, Quantum neuromorphic computing with reservoir computing networks, *Adv. Quantum Technol.* **4**, 2100053 (2021).
 - [9] P. Mujal, R. Martínez-Peña, J. Nokkala, J. García-Bení, G. L. Giorgi, M. C. Soriano, and R. Zambrini, Opportunities in quantum reservoir computing and extreme learning machines, *Adv. Quantum Technol.* **4**, 2100027 (2021).
 - [10] K. Nakajima, K. Fujii, M. Negoro, K. Mitarai, and M. Kitagawa, Boosting computational power through spatial multiplexing in quantum reservoir computing, *Phys. Rev. Appl.* **11**, 034021 (2019).
 - [11] R. Martínez-Peña, G. L. Giorgi, J. Nokkala, M. C. Soriano, and R. Zambrini, Dynamical phase transitions in quantum reservoir computing, *Phys. Rev. Lett.* **127**, 100502 (2021).
 - [12] P. Mujal, J. Nokkala, R. Martínez-Peña, G. L. Giorgi, M. C. Soriano, and R. Zambrini, Analytical evidence of nonlinearity in qubits and continuous-variable quantum reservoir computing, *J. Phys.: Complex.* **2**, 045008 (2021).
 - [13] P. Mujal, R. Martínez-Peña, G. L. Giorgi, M. C. Soriano, and R. Zambrini, Time-series quantum reservoir computing with weak and projective measurements, *Npj Quantum Inf.* **9**, 16 (2023).
 - [14] J. Chen, H. I. Nurdin, and N. Yamamoto, Temporal information processing on noisy quantum computers, *Phys. Rev. Appl.* **14**, 024065 (2020).
 - [15] T. Monomi, W. Setoyama, and Y. Hasegawa, Feedback-enhanced quantum reservoir computing with weak measurements (2025), arXiv:2503.17939 [quant-ph].
 - [16] U. Hübner, N. B. Abraham, and C. O. Weiss, Dimensions and entropies of chaotic intensity pulsations in a single-mode far-infrared nh_3 laser, *Phys. Rev. A* **40**, 6354 (1989).
 - [17] A. Weigend and N. Gershenfeld, Results of the time series prediction competition at the santa fe institute, in *IEEE International Conference on Neural Networks* (1993) pp. 1786–1793 vol.3.
 - [18] A. D. Córcoles, M. Takita, K. Inoue, S. Lekuch, Z. K. Mineev, J. M. Chow, and J. M. Gambetta, Exploiting dynamic quantum circuits in a quantum algorithm with superconducting qubits, *Phys. Rev. Lett.* **127**, 100501 (2021).
 - [19] J. M. Pino, J. M. Dreiling, C. Figgatt, J. P. Gaebler, S. A. Moses, M. Allman, C. Baldwin, M. Foss-Feig, D. Hayes, K. Mayer, *et al.*, Demonstration of the trapped-ion quantum ccd computer architecture, *Nature* **592**, 209 (2021).
 - [20] M. DeCross, E. Chertkov, M. Kohagen, and M. Foss-Feig, Qubit-reuse compilation with mid-circuit measurement and reset, *Phys. Rev. X* **13**, 041057 (2023).
 - [21] Y. Yu, K. Yan, D. Biswas, V. N. Zhang, B. Harraz, C. Noel, C. Monroe, and A. Kozhanov, In situ midcircuit qubit measurement and reset in a single-species trapped-ion quantum computing system, *Phys. Rev. Res.* **7**, 043355 (2025).