

A hidden bottleneck in classical and quantum linear reservoir computing¹

Johannes Nokkala^{1,2}, Federico Centrone^{3,4}, Francesco Arzani⁵

¹Department of Physics and Astronomy, University of Turku, Finland

²Turku Collegium for Science, Medicine and Technology, University of Turku, Finland

³ICFO–Institut de Ciències Fòniques, The Barcelona Institute of Science and Technology, Spain

⁴Universidad de Buenos Aires, Facultad de Ciencias Exactas y Naturales, Argentina

⁵QAT team, DIENS, École Normale Supérieure, PSL University, CNRS, INRIA, Paris, France

Summary. We identify a hidden bottleneck in the information processing capacity (IPC) of linear reservoir computers. Since Gaussian transformations act linearly on covariance matrices, this bound applies directly to Gaussian continuous-variable (CV) quantum reservoirs. We show that experimentally accessible single-photon operations can overcome it, and that the resulting excess capacity provides an operational witness of non-Gaussian processing.

1 Background and motivation

Reservoir computing circumvents the difficulties of training recurrent neural networks by keeping the dynamical system (the reservoir) fixed and training only a linear readout [1]. This paradigm is especially attractive for quantum platforms, where engineering controllable nonlinear dynamics is a central challenge. However, the reliance on a linear readout raises a fundamental question: how much of the computational power originates from the reservoir itself, and how much is already present in the input encoding?

This question is particularly pressing for continuous-variable quantum reservoir computing (QRC) with photonic platforms [2, 3], where Gaussian state preparation and Gaussian dynamics are the most experimentally accessible ingredients. Since Gaussian transformations act linearly on quadratures and covariance matrices, the computational power available to a purely Gaussian reservoir is structurally constrained.

2 The linear-reservoir bottleneck

The goal of reservoir computing is to approximate temporal input–output maps: given an input time series s_t , the reservoir should reproduce a desired target function of the recent input history. The IPC [4] quantifies how large the class of such functions that a given reservoir can approximate is, independently of any single task. Given an i.i.d. input time series s_t , the IPC is resolved into contributions from Legendre basis functions $P_d(s_{t-\tau})$ labeled by nonlinearity degree d and delay τ . The delay-resolved capacity $\text{IPC}^{(\tau)} = \sum_{d \geq 0} C_{P_d^\tau}$ measures the expressive power available at a specific delay.

We consider reservoirs whose observables $x_t \in \mathbb{R}^m$ admit a finite-dimensional linear update

$$x_t = \mathbf{A} x_{t-1} + \mathbf{B} g_t, \tag{1}$$

¹Preprint available at <https://arxiv.org/abs/2605.29071>

where $g_t \in \mathbb{R}^p$ is the preprocessed input, and the output is formed by linear readout with bias, $o_t = \mathbf{w}^\top x_t + w_0$.

Theorem 1. *Consider a reservoir of the form (1) with linear readout and bias.*

- (A) *If preprocessing has memory but the reservoir is memoryless ($\mathbf{A} = \mathbf{0}$), the total IPC after the reservoir is upper bounded by the IPC of the preprocessing features alone.*
- (B) *If preprocessing is memoryless but the reservoir has linear memory, the delay-resolved IPC at every fixed delay τ is upper bounded by the IPC of the preprocessing features alone.*

The proof relies on two observations. In case (A), a linear map cannot increase the rank of the feature matrix. In case (B), unrolling Eq. (1) expresses each observable as a sum of terms depending on one delay at a time. Since basis functions at different delays are orthogonal for i.i.d. inputs, only the matching-delay term can contribute to the reconstruction of a target at a given delay.

This bottleneck is *hidden* in aggregate measures: contributions from different delays can accumulate, so the total IPC may remain large even when the expressive power at each individual delay is sharply constrained.

3 Gaussian bound for CV quantum reservoirs

In CV QRC, the readout is commonly formed from elements of the covariance matrix of measured quadratures. Gaussian transformations act on the covariance matrix as $\sigma \mapsto S\sigma S^\top$, which after vectorization is a linear map. Covariance-based Gaussian reservoirs therefore satisfy Eq. (1) and fall within the scope of Theorem 1.

Concretely, for a reservoir of N modes with memoryless input encoding, the Gaussian limit on the IPC at any fixed delay scales linearly with N (set by the number of independent encoding parameters), whereas the number of available covariance-matrix observables scales as $N(N+1)/2$. The gap between these two scalings defines the region that is forbidden to Gaussian reservoirs but may become accessible through non-Gaussian resources.

4 Breaking the bound with single-photon operations

Motivated by recent experimental demonstrations of mode-selective single-photon subtraction [5] and addition [6] in multimode quantum fields, as well as the first experimental realization of continuous-variable quantum reservoir computing [3], we consider a photonic reservoir scheme in which an input time series modulates squeezed-vacuum states, correlations are created by a fixed multimode Gaussian transformation, and non-Gaussianity is introduced through conditional single-photon subtraction or addition. The readout is formed from homodyne statistics, keeping the scheme close to existing Gaussian photonic reservoirs. The non-Gaussian observables are evaluated exactly, without Hilbert-space truncation, using Wick's theorem applied to creation and annihilation operators acting on the underlying Gaussian state [7].

Our numerical results demonstrate the following:

- In the memoryless setting (quantum extreme learning machine), single-photon operations generate capacities that exceed the Gaussian bound, entering the region forbidden to all Gaussian reservoirs. The excess grows with additional non-Gaussian operations.




Figure 1: Breaking the Gaussian limit in the memoryless case. Left: total capacity vs. number of modes N . The shaded region is inaccessible to Gaussian reservoirs but becomes reachable with single-photon operations. Right: capacity distributed across polynomial degrees at $N = 10$, showing that saturation requires support on increasingly high-degree functions. Averaged over 80 realizations.

- Introducing classical memory dramatically improves efficiency: the non-Gaussian reservoir approaches the absolute capacity limit much faster, because expressive power can be assembled from products of lower-degree functions across multiple delays rather than requiring high-degree functions at a single delay.
- The excess capacity degrades gracefully under finite-ensemble homodyne sampling modeled by the Wishart distribution, remaining visible at experimentally realistic ensemble sizes.

Corollary 1. *Under the assumptions of Theorem 1, any observed delay-resolved capacity exceeding the Gaussian bound rules out a Gaussian explanation of the effective reservoir transformation.*

This provides an operational witness of non-Gaussian processing extracted directly from input–output data, without requiring state tomography.

5 Conclusions and outlook

We have identified a structural limitation of linear, and in particular Gaussian, reservoir transformations for information processing. The bottleneck is delay-resolved and can be invisible in aggregate capacity measures. Experimentally accessible single-photon operations overcome this limitation, and the resulting excess capacity serves as an operational witness of non-Gaussian processing. These results clarify which resources are genuinely responsible for enhanced expressive power in CV QRC, and suggest that distributing non-Gaussian processing across multiple delays is far more effective than concentrating it into a memoryless architecture.

References

- [1] Mantas Lukoševičius and Herbert Jaeger. Survey: Reservoir computing approaches to recurrent neural network training. *Comput. Sci. Rev.*, 3(3):127–149, August 2009.
- [2] Johannes Nokkala, Rodrigo Martínez-Peña, Gian Luca Giorgi, Valentina Parigi, Miguel C Soriano, and Roberta Zambrini. Gaussian states of continuous-variable quantum systems provide universal and versatile reservoir computing. *Communications Physics*, 4(1):53, 2021.
- [3] Iris Paparelle, Johan Henaff, Jorge Garcia-Beni, Emilie Gillet, Daniel Montesinos, Gian Luca Giorgi, Miguel C Soriano, Roberta Zambrini, and Valentina Parigi. Experimental memory control in continuous-variable optical quantum reservoir computing. *Nature Photonics*, 20:413–420, 2026.
- [4] Joni Dambre, David Verstraeten, Benjamin Schrauwen, and Serge Massar. Information processing capacity of dynamical systems. *Scientific reports*, 2(1):514, 2012.
- [5] Young-Sik Ra, Adrien Dufour, Mattia Walschaers, Clément Jacquard, Thibault Michel, Claude Fabre, and Nicolas Treps. Non-Gaussian quantum states of a multimode light field. *Nature Physics*, 16(2):144–147, February 2020. Publisher: Nature Publishing Group.
- [6] Ganaël Roeland, Srinivasan Kaali, Victor Roman Rodriguez, Nicolas Treps, and Valentina Parigi. Mode-selective single-photon addition to a multimode quantum field. *New Journal of Physics*, 24(4):043031, 2022.
- [7] Paolo Stornati, Antonio Acin, Ulysse Chabaud, Alexandre Dauphin, Valentina Parigi, and Federico Centrone. Variational quantum simulation using non-gaussian continuous-variable systems. *Physical Review Research*, 6(4):043212, 2024.