

Rydberg-Atom Quantum Reservoir Computing for Insurance Time-Series Forecasting

Nathaniel James Pacey^{1,2}, George Woodman¹

¹AXA Group Operations, Paris, France, ²EPFL, Lausanne, Switzerland

Quantum machine learning (QML) can generate non-classical feature spaces through superposition, interference, entanglement, and many-body correlations [1, 2]. This is promising for forecasting from limited, nonlinear, and event-driven data. Here, we use daily road-accident records as auto-insurance claim proxies, since collisions, casualties, and weather-related incidents drive motor losses. These sparse, seasonal, nonlinear series are difficult to forecast due to reporting effects, mobility, weather, and external events. Many gate-based QML methods use parameterized quantum circuits (PQCs), which on current NISQ hardware can suffer from barren plateaus, noise-induced trainability loss, and sensitivity to circuit depth [3, 4, 5]. Quantum reservoir computing (QRC) provides a more hardware-compatible alternative by fixing the quantum dynamics as a computational reservoir and training only a classical readout [6]. This allows QRC to turn many-body dynamics, memory effects, and responses from noise or dissipation into useful computational features, without requiring deep, idealized circuits. In this work, we extend the Rydberg-array QRC framework of Liu et al. [7] to insurance-relevant forecasting. We optimize Rydberg QRC on synthetic benchmarks, including NARMA- n tasks and damped harmonic oscillators, then test the optimized models on next-step prediction for daily auto-risk proxy series. We compare several Rydberg QRC configurations against classical forecasting baselines, including echo state networks (ESNs), ARIMA and ETS and prepare an optimized setup for execution on QuEra’s Aquila analog neutral-atom processor [8]. Running this model on real hardware will test whether analog QRC is practical on current noisy devices for sparse, event-driven insurance-related forecasting.

Our reservoir is an array of N two-level atoms governed by the input-dependent Rydberg Hamiltonian

$$H(u_t) = \sum_i \frac{\Omega_i(u_t)}{2} X_i - \sum_i \Delta_i(u_t) n_i + \sum_{i < j} \frac{C_6}{\|r_i - r_j\|^6} n_i n_j, \quad n_i = \frac{1 - Z_i}{2}. \quad (1)$$

The scalar time-series input is encoded through drive and detuning modulation, $\Omega_i(u_t) = \Omega_{0,i} + a_i u_t$ and $\Delta_i(u_t) = \Delta_{0,i} + b_i u_t$. The reservoir evolves for a fixed interval, with optional shot noise, dephasing, relaxation, and hardware sampling models. We compare multi-step QRC (MS-QRC), where the quantum state carries memory between inputs, and single-step QRC (SS-QRC), where each input or input window is independently encoded. Classical features are formed from simulated or measured observables, including one- and two-body expectation values in simulation and occupation/correlation features from Aquila-compatible bitstrings. A time-stacked feature vector $\phi_t = [1, z_t, z_{t-1}, \dots, z_{t-k+1}]$ is mapped to the forecast by ridge regression,

$$\hat{y}_{t+h} = W_{\text{out}} \phi_t, \quad W_{\text{out}} = (\Phi^T \Phi + \lambda I)^{-1} \Phi^T Y. \quad (2)$$

Our empirical study uses STATS19 road-safety data, which include accident, vehicle, casualty, date, location, road-condition, weather, and severity fields (collected from 2020-01-01 to 2024-12-31 total of 1827 days) [12]. From these records, we construct three daily forecasting targets: collision count, casualty count, and collisions on wet, icy, snowy, flooded, or otherwise adverse road surfaces. Each target is evaluated using a chronological 70/30 train–test split. For QRC, we sweep qubit number, preprocessing choice: identity scaling or log1p with min–max scaling, and washout, selecting washout 200 for the reported results. The ESN baseline uses a 200-unit leaky-integrator reservoir, with hyperparameters selected from 5000 random NARMA10 configurations over literature-guided ranges: $\rho \in [0.8, 1.0]$, $\alpha \in [0.1, 0.5]$, $s_{\text{in}} \in [0.1, 1.0]$, and washout 50–100 steps [9, 10, 11]. The best ESN configurations clustered near $\rho = 0.9$, $\alpha = 0.3$, $s_{\text{in}} = 1.0$, and washout 60, which were used for the STATS19 comparison.

Task	ETS	ARIMA	Best ESN	Best SS-QRC	Best MS-QRC
Daily collision count	-0.083	-0.018	0.317 (identity)	0.483 (7, identity)	0.545 (6, identity)
Daily casualty count	-0.074	-0.024	0.255 (log1p)	0.388 (7, log1p)	0.461 (6, identity)
Wet/icy-road collision count	-0.366	-0.050	0.413 (log1p)	0.458 (6, log1p)	0.464 (7, identity)

All entries report R_{test}^2 . ESN entries list the best preprocessing choice; QRC entries list the best configuration as (qubits, preprocessing).

These preliminary results show that small simulated Rydberg reservoirs can outperform a tuned ESN baseline on noisy, sparse, insurance-relevant forecasting tasks. MS-QRC outperforms the ESN baseline on all three targets, with absolute R_{test}^2 gains ranging from +0.052 to +0.228, indicating that limited atom numbers can already generate useful nonlinear temporal features for actuarial data. Continued sweeps over larger reservoirs and comparisons with classical models will test whether richer many-body dynamics and analog feature generation further improve performance. Since QuEra’s Aquila device supports up to 256 atoms [8], hardware-scale reservoirs may access substantially larger dynamical feature spaces and provide a practical route to probe potential QML advantage in the NISQ era.

References

- [1] J. Biamonte, P. Wittek, N. Pancotti, P. Rebentrost, N. Wiebe, and S. Lloyd, “Quantum machine learning,” *Nature*, vol. 549, pp. 195–202, 2017.
- [2] M. Cerezo *et al.*, “Challenges and opportunities in quantum machine learning,” *Nature Computational Science*, vol. 2, pp. 567–576, 2022.
- [3] J. R. McClean, S. Boixo, V. N. Smelyanskiy, R. Babbush, and H. Neven, “Barren plateaus in quantum neural network training landscapes,” *Nature Communications*, vol. 9, Art. no. 4812, 2018.
- [4] S. Wang *et al.*, “Noise-induced barren plateaus in variational quantum algorithms,” *Nature Communications*, vol. 12, Art. no. 6961, 2021.
- [5] J. Preskill, “Quantum computing in the NISQ era and beyond,” *Quantum*, vol. 2, Art. no. 79, 2018.
- [6] K. Fujii and K. Nakajima, “Harnessing disordered-ensemble quantum dynamics for machine learning,” *Physical Review Applied*, vol. 8, Art. no. 024030, 2017.
- [7] D.-S. Liu *et al.*, “Practical quantum reservoir computing in Rydberg atom arrays,” arXiv:2602.00610, 2026.
- [8] Amazon Braket, “QuEra quantum computers: Aquila processor.” [Online]. Available: <https://aws.amazon.com/braket/quantum-computers/quera/> Accessed: Jul. 8, 2026.
- [9] H. Jaeger, “The ‘echo state’ approach to analysing and training recurrent neural networks,” GMD Report 148, German National Research Center for Information Technology, 2001.
- [10] H. Jaeger, M. Lukoševicius, D. Popovici, and U. Siewert, “Optimization and applications of echo state networks with leaky-integrator neurons,” *Neural Networks*, vol. 20, no. 3, pp. 335–352, 2007.
- [11] M. Lukoševicius, “A practical guide to applying echo state networks,” in *Neural Networks: Tricks of the Trade*, 2nd ed., G. Montavon, G. B. Orr, and K.-R. Müller, Eds. Berlin, Germany: Springer, 2012, pp. 659–686.
- [12] UK Department for Transport, “Road safety open data.” [Online]. Available: <https://www.gov.uk/government/statistical-data-sets/road-safety-open-data> Accessed: Jul. 8, 2026.