

A quantum-inspired spiking variational autoencoder for robust anomaly detection at the LHC

Dean Brand^{1,3}, Donovan Slabbert^{1,3}, Chanelle Matadah Manfouo¹, Francesco Petruccione^{2,3}

¹*Department of Physics, Stellenbosch University, Stellenbosch, South Africa*

²*School for Data Science and Computational Thinking, Stellenbosch University, Stellenbosch, South Africa*

³*National Institute for Theoretical and Computational Sciences (NITheCS), South Africa*

Model-agnostic searches for new physics at the LHC train on Standard-Model background and flag atypical events, an unsupervised anomaly-detection task [1, 2] for which quantum [3, 4] and neuromorphic detectors are increasingly proposed. Whether these detectors match strong classical baselines, and whether their dynamics add anything beyond an ordinary spiking network, is still untested. We develop the quantum-inspired spiking variational autoencoder (**QSVAE**), the first generative model built on the quantum-inspired leaky integrate-and-fire (**QLIF**) neuron. We recast the QLIF neuron from the measurement-based quantum circuit of Brand and Petruccione [5] into a differentiable component and place it inside the fully spiking variational autoencoder (**FSVAE**) of Kamata *et al.* [6]. Its bounded, periodic $\sin^2$ -rotation membrane replaces the classical neuron at equal weight count, so the FSVAE and QSVAE form an architecture-matched pair that isolates the effect of the neuron.

We build the benchmark on the di-jet latent space of Belis *et al.* [3], where a one-class detector trains on QCD background and is tested against Randall–Sundrum gravitons and an $A \rightarrow HZ \rightarrow ZZZ$ scalar. We place fourteen detectors on this representation: model-free statistics; classical shallow (k -nearest-neighbour, Mahalanobis, Isolation Forest) and deep (autoencoder, Deep SVDD) detectors; the spiking FSVAE and QSVAE; and a quantum family. We reproduce the published quantum suite [3] of entangled and non-entangled quantum kernels, quantum k -means and k -medians, and a classical RBF reference, recovering the kernel’s advantage over that reference (0.855 against 0.755 at $\ell = 16$). We then add three quantum autoencoders [4]: a persistent variant (**pQAE**), a bottleneck variant (**bQAE**), and a **hybrid** that pairs a quantum encoder with a classical decoder. On clean data the simplest methods win. Ranked by the area under the ROC curve (AUROC), the k -nearest-neighbour distance leads at every latent dimension (Overall 0.891 at $\ell = 16$), the FSVAE ranks among the strongest learned models (0.877), and the QSVAE sits mid-pack (0.827). Among the quantum autoencoders the hybrid is the strongest (0.877, non-entangled) and matches the spiking autoencoders, and the fully quantum pQAE and bQAE score lower, held back by their quantum decoders.

The QSVAE earns its place through robustness to input perturbation. Under a fair, energy-matched, geometry-respecting perturbation ($\sigma = 0.3$) it retains 0.95 of its clean accuracy, and twelve classical-spiking controls that vary firing rate, saturation shape, and regularisation all fall in a single 0.74 to 0.85 band. A firing-matched ablation identifies the periodic $\sin^2$ dynamics as the common source of this robustness and of the QSVAE’s clean-accuracy cost, since bounded-classical neurons matched to its firing rate reproduce neither, and the robustness tracks a 3.5-fold reduction in the response of the anomaly score to perturbation. Two controls qualify the result: the robustness inversion seen under isotropic noise is an artefact of the noise model, and we claim no quantum advantage. The study delivers a fair cross-family benchmark and a controlled account of how quantum-inspired spiking dynamics decouple noise robustness from model capacity.

References

- [1] G. Kasieczka *et al.*, “The LHC Olympics 2020: A community challenge for anomaly detection in high energy physics,” *Rep. Prog. Phys.* **84**, 124201 (2021).
- [2] T. Aarrestad *et al.*, “The Dark Machines anomaly score challenge,” *SciPost Phys.* **12**, 043 (2022).
- [3] V. Belis *et al.*, “Quantum anomaly detection in the latent space of proton collision events at the LHC,” *Commun. Phys.* **7**, 334 (2024).
- [4] V. S. Ngairangbam, M. Spannowsky, and M. Takeuchi, “Anomaly detection in high-energy physics using a quantum autoencoder,” *Phys. Rev. D* **105**, 095004 (2022).
- [5] D. Brand and F. Petruccione, “A quantum leaky integrate-and-fire spiking neuron and network,” *npj Quantum Inf.* **10**, 125 (2024).
- [6] H. Kamata, Y. Mukuta, and T. Harada, “Fully Spiking Variational Autoencoder,” *Proc. AAAI Conf. Artif. Intell.* **36**, 7059–7067 (2022).