

Hybrid Quantum-Classical Logistic Regression for Calibrated Classification of Pulsar Candidates

Chanelle Matadah Manfouo^{1,2,3}, Donovan Slabbert^{1,3}, Prince Koree Osei^{4,2}, Francesco Petruccione^{3,5,1}

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

²African Institute for Mathematical Sciences Research and Innovation Centre, Kigali, Rwanda

³National Institute for Theoretical and Computational Sciences, Stellenbosch, South Africa

⁴African Institute for Mathematical Sciences, Accra, Ghana

⁵School of Data Science and Computational Thinking, Stellenbosch University, South Africa

Modern radio surveys produce large numbers of candidate signals that must be screened to identify new pulsars. This is a rare-event classification problem, since most candidates correspond to noise or radio-frequency interference rather than true pulsars. The HTRU-2 dataset of Lyon et al. [1] is a standard benchmark for this task and contains eight statistical features extracted from pulsar diagnostic plots, with only about nine percent pulsar candidates. In this setting, candidate ranking is as important as binary classification. A model should therefore provide probability estimates that are not only discriminative, but also well calibrated for telescope follow-up decisions.

Classical machine learning methods have been widely used for pulsar candidate selection, and recent quantum machine learning studies have explored single-qubit quantum neural networks, quantum support vector machines, quantum convolutional neural networks, and shallow variational circuits for related pulsar classification tasks [1–4]. These studies mainly assess discrimination through metrics such as the accuracy, the precision, the recall, the Area Under the ROC curve (AUROC), and the Area Under the Precision-Recall Curve (AUPR). Less attention has been given to probability calibration, although candidate follow-up depends on whether predicted probabilities agree with observed pulsar frequencies.

We study hybrid quantum-classical logistic regression for calibrated pulsar candidate classification [10]. The model uses a parameterised quantum circuit as a trainable feature map and passes Pauli- Z expectation values to a classical logistic layer. This preserves the probabilistic interpretation of logistic regression while allowing the input features to be transformed by a quantum circuit. Three data-encoding strategies are compared in the same framework: angle encoding and amplitude encoding, following standard quantum machine learning constructions [5], and data re-uploading, where the input features are introduced repeatedly across circuit layers [6]. The quantum logistic models are trained with binary cross-entropy and compared with classical baselines, including, logistic regression, support vector machine (SVM), XGBoost, and a quantum SVM as a reference model.

The evaluation combines discrimination, low-false-positive-rate recovery, calibration, and runtime. Since pulsars are the minority class, AUROC and AUPR are reported together with Recall at low (1%, 5%) false-positive rate, which measures pulsar recovery under a strict false-positive-rate constraint. Calibration is assessed using the Brier score, expected calibration error, reliability diagrams, and Murphy decomposition [7–9]. A paired-seed protocol is used so that all models are evaluated on matched stratified train-test splits and matched random seeds.

Across the tested training dataset sizes and circuit depths, angle-encoded quantum logistic regression gives the strongest performance among the quantum logistic regression variants. It remains close to the best classical baseline in AUPR, AUROC, and Recall at low false-positive rate, while giving the lowest calibration error among the compared models. Classical logistic regression remains strong in discrimination, but its probability estimates are less well aligned with empirical class frequencies in this benchmark. Murphy decomposition links the lower calibration error of the angle-encoded model to low reliability error and stable resolution, so the model combines reliable probability estimates with good separation between the classes.

The three encoding strategies show different behaviour across circuit depths and training-set sizes. Angle encoding is the most stable strategy in the tested settings, and increasing circuit depth does not consistently improve its performance. Data re-uploading is competitive at shallow depth but loses discrimination and resolution at larger depth in the present multi-qubit implementation. Amplitude encoding has the lowest computational cost among the three quantum logistic models, but remains weaker in predictive performance. The results identify shallow angle-encoded quantum logistic regression as the best-balanced quantum logistic model for this rare-event screening task, and emphasise the role of calibration and operational metrics when quantum machine learning predictions are intended to support scientific decision-making.

References

- [1] R. J. Lyon, B. W. Stappers, S. Cooper, J. M. Brooke, and J. D. Knowles, “Fifty years of pulsar candidate selection: from simple filters to a new principled real-time classification approach,” *Monthly Notices of the Royal Astronomical Society* 459, 1104–1123 (2016).
- [2] M. Kordzanganeh, A. Utting, and A. Scaife, “Quantum machine learning for radio astronomy,” arXiv:2112.02655 (2021).
- [3] D. Slabbert, M. Lourens, and F. Petruccione, “Pulsar classification: comparing quantum convolutional neural networks and quantum support vector machines,” *Quantum Machine Intelligence* 6, 56 (2024).
- [4] S. McFarthing, A. Pillay, I. Sinayskiy, and F. Petruccione, “Classical ensembles of single-qubit quantum variational circuits for classification,” *Quantum Machine Intelligence* 6, 81 (2024).
- [5] M. Schuld and F. Petruccione, *Machine Learning with Quantum Computers*, Springer (2021).
- [6] A. Pérez-Salinas, A. Cervera-Lierta, E. Gil-Fuster, and J. I. Latorre, “Data re-uploading for a universal quantum classifier,” *Quantum* 4, 226 (2020).
- [7] W. B. Glenn et al., “Verification of forecasts expressed in terms of probability,” *Monthly Weather Review* 78, 1–3 (1950).
- [8] A. H. Murphy, “A new vector partition of the probability score,” *Journal of Applied Meteorology* 12, 595–600 (1973).
- [9] C. Guo, G. Pleiss, Y. Sun, and K. Q. Weinberger, “On calibration of modern neural networks,” in *Proceedings of the International Conference on Machine Learning* (2017).
- [10] C. M. Manfouo, D. Slabbert, P. K. Osei, and F. Petruccione, “Hybrid quantum-classical logistic regression for calibrated classification of pulsar candidates,” arXiv:2605.07866 (2026).