Qeios

Peer Review

Review of: "Hybrid Quantum Neural Networks with Amplitude Encoding: Advancing Recovery Rate Predictions"

Ishan Shivansh Bangroo[1]

1. IEEE Computer Society, Washington D.C., United States

The manuscript *Hybrid Quantum Neural Networks with Amplitude Encoding: Advancing Recovery Rate Predictions* proposes a hybrid quantum-classical neural network for recovery rate prediction. A parameterised quantum circuit replaces one hidden layer of a classical neural network and uses amplitude encoding to map 256 input features towards a smaller number of qubits. The model is evaluated on a proprietary global dataset with 1725 defaulted bonds and compared with a classical feed-forward neural network and a quantum model using angle encoding. Models are tuned to have broadly comparable numbers of trainable parameters and are evaluated with four-fold cross-validation using root mean squared error (RMSE). The amplitude encoding model achieves the lowest average test RMSE and the smallest standard deviation among the three models. The research is well motivated from the perspective of credit risk, representing a useful contribution to the growing literature on hybrid quantum models in finance.

The paper explains why accurate recovery rate prediction matters for pricing and risk management, why simple fixed recovery assumptions are inadequate, and how the limited number of defaults relative to feature dimension motivates the exploration of regularised or constrained models. The experimental design is disciplined; the authors use k-fold cross-validation, reporting RMSE together with standard deviations across folds, and include auxiliary experiments in the appendices which vary the number of hidden units in the classical model and the number of qubits in the angle encoding model. These additional experiments provide useful insight into overfitting in larger classical networks and saturation behaviour in the angle encoding model. However, there are several areas where the technical content could be strengthened:

1. Clarify what is meant by *quantum advantage* in the context; all experiments are run on a classical state vector simulator without noise, and the main benefit observed is an improvement in predictive accuracy and stability for a particular hybrid architecture. It would help to state clearly that the results show modelling advantage under simulation rather than a computational speed-up on quantum hardware. A short discussion of the depth and gate count of the amplitude-preparation circuit and how this scales as the number of features grows would make the comparison more transparent.
2. We should reconsider the use of test folds for selecting the best epoch; *Table* 4 reports the best RMSE on the test data across 100 epochs for each model. This effectively uses the test data for model selection, which can bias RMSE downward. A standard alternative would be to select the stopping epoch on a validation subset of the training data and then report the RMSE on the held-out test data. At minimum, the current procedure should be described explicitly as a limitation; ideally, results would be recomputed with a separate validation step to confirm that the performance ordering remains unchanged.
3. Strengthening and contextualising the baseline models with the feed-forward neural network baseline and the angle encoding quantum baseline are reasonable and are explored well in the appendices; however, recovery rate prediction is often tackled with tree-based models or other classical non-linear methods. Including at least one strong non-neural classical baseline (like gradient-boosted trees) or discussing why such models were not included would help readers judge the size and relevance of the improvement achieved by the hybrid quantum model.
4. Expand the error analysis and diagnostics beyond average RMSE and its standard deviation; additional metrics like mean absolute error or the coefficient of determination could provide a fuller picture. Plots of predicted versus observed recovery rates, possibly stratified by rating, seniority, or calendar time, would show whether the hybrid quantum model provides particular benefit in certain regions of the recovery distribution.
5. Tighten notation and correct minor presentation issues; there are a few typographical errors and formatting glitches, such as in the expression for qubit scaling and in some appendix headings. A careful proofreading to standardise notation, fix small errors, and ensure that all symbols are defined would improve readability, especially for readers whose primary background is in finance rather than quantum information.
6. Clarify the extent of reproducibility because the dataset is proprietary and cannot be shared in full; it would be helpful to state exactly what could be reproduced by other researchers. The authors

might consider releasing full model definitions, training scripts, and hyperparameter settings, possibly together with a synthetic or public surrogate dataset on which the architecture can be demonstrated. This would enhance the impact of the work in the broader community.

These are some refinements rather than fundamental objections. The core idea and implementation are good, and the results are promising. Addressing the points above would strengthen the statistical validity of the conclusions, make the comparison with classical models more informative, and clarify how the proposed hybrid quantum architecture might be used and extended in future credit risk applications.

Declarations

Potential competing interests: No potential competing interests to declare.