

Hybrid Classical–Quantum Learning for Job Failure Prediction in Distributed Cloud Systems

Muhammad Afaq*, Zeeshan Kaleem**, Muhammad Ibrahim[†], Redowan Mahmud[‡]

*Computer Engineering Department and IRC for Intelligent Secure Systems, KFUPM, Saudi Arabia

**Computer Engineering Department and IRC for Smart Mobility and Logistics, KFUPM, Saudi Arabia

[†]Department of Information Technology, University of Haripur, Pakistan

[‡]School of Electrical Engineering, Computing and Mathematical Sciences, Curtin University, Australia

{muhammad.afaq, zeeshan.kaleem}@kfupm.edu.sa, ibrahimmayar@uoh.edu.pk, mdredowan.mahmud@curtin.edu.au

Abstract—Accurate job failure predictions in distributed cloud environments are vital for improving reliability, scheduling efficiency, and cost-effectiveness. Moreover, the dynamic variability and scale of production workloads make failure prediction a challenge to overcome. Although classical machine learning (ML) models, particularly Random Forests (RF) provide strong predictive performance, they can be computationally demanding under frequent retraining. Contrarily, Quantum Machine Learning (QML) methods leverage quantum-enhanced kernels to capture complex patterns, but standalone quantum models continue to have hardware limitations related and depict significant computational overhead. In this paper, we examine a hybrid classical-quantum ML framework for job failure prediction in large-scale distributed systems. A Hybrid Quantum Support Vector Machine (QSVM) is employed, blending classical RBF and quantum kernels, and benchmarked against RF. Experimental results show that RF attains high predictive accuracy ($\text{Acc} \approx 0.98 - 0.99$, $\text{ROC AUC} \approx 0.99$, $\text{AP} \approx 0.98$), while hybrid QSVM achieves moderate accuracy ($\approx 0.80 - 0.82$), but they exhibit distinct timing behavior. While the quantum kernel evaluation is computationally expensive, the fitting is faster when the kernels are precomputed. This aspect offers potential advantages in cloud environments where rapid retraining is needed under evolving workloads. Furthermore, our results emphasize that classical and quantum models are not competing but rather complementing each other. While classical models offer high accuracy, the hybrid quantum approaches provide efficiency trade-offs for large-scale, time-sensitive prediction tasks.

Index Terms—job failure prediction, distributed cloud systems, classical machine learning, quantum machine learning, quantum support vector machine

I. INTRODUCTION

Cloud computing has become the backbone of modern digital infrastructure, allowing organizations to run baseload applications efficiently, flexibly, and cost-effectively on a larger scale. [1], [2]. However, the increasing complexity of cloud datacenters, diverse workloads, and dynamic scheduling policies present notable challenges, especially in precisely predicting job behaviors and avoiding failures. Job failure prediction is vital, as failures have a direct impact on service-level objectives (SLOs), energy consumption, resource utilization, and system reliability [3], [4]. Therefore, improving predictive models for failure detection portrays a crucial research objective in cloud operations, with direct significance for enhancing quality of service (QoS), efficiency, and resilience.

Classical ML models, such as logistic regression, RF, and support vector machines (SVMs) have been widely used in cloud research for tasks including workload characterization, anomaly detection, and job scheduling [5], [6]. While these models perform adequately, they are generally unable to accurately model nonlinear, complex, and high-dimensional interactions that exist in large-scale distributed systems, especially when datacenter states fluctuate or workloads evolve rapidly [7]. Recently, QML has emerged as a promising research venue, leveraging particular characteristics of quantum computing, such as entanglement and superposition to capture complex correlations in data that may be neglected by purely classical approaches [8], [9].

Despite the potential of QML, practical usability currently is constrained by the limitations of quantum hardware, including limited qubit count, noise and coherence times [10]. These limitations have fostered hybrid quantum–classical approaches, where classical ML provides stability and efficiency, while quantum kernels enhance feature representations and enrich generalization [11]–[13]. Although QML has been explored in several fields, there remains a notable research gap in its study within the area of cloud job failure prediction.

This paper addresses this gap by proposing a hybrid classical–quantum ML framework for job failure prediction in distributed cloud systems. The objectives of this work are twofold: first, to assess the practicability and effectiveness of a hybrid classical–quantum model for job failure prediction in distributed cloud clusters; and second, to provide empirical evidence that quantum-enhanced techniques can complement classical methods by offering accuracy–efficiency trade-offs in large-scale prediction tasks. Moreover, this paper aims to exhibit how a hybrid learning framework can improve predictive capabilities in cloud environments, setting the stage for future developments as quantum computing technology progresses.

The remainder of this paper is organized as follows: Section II presents the system architecture and methodology, detailing the hybrid pipeline, including classical preprocessing, quantum kernel mapping, and model integration. Section III provides a comprehensive performance evaluation using accuracy, F1 score, ROC AUC, average precision, and training time, supported by comparative visualizations. It also includes




Fig. 1. System architecture of the proposed hybrid classical–quantum framework for job failure prediction. The pipeline includes classical preprocessing, quantum feature embedding using ZZFeatureMap, and classification via Random Forest and Hybrid QSVM.

a discussion on trade-offs and the complementary role of quantum components. Finally, Section IV concludes the paper and outlines directions for future research in quantum-enhanced cloud intelligence.

II. SYSTEM ARCHITECTURE AND METHODOLOGY

The proposed system architecture, as illustrated in Fig. 1, follows a layered, modular design wherein classical and quantum components are integrated in a complementary fashion to address the job failure prediction problem in distributed cloud environments. The overall objective is to accurately predict whether a submitted job will complete successfully or fail, based on features such as scheduling class, requested CPU, and requested memory.

Let the dataset be represented as a collection of job instances:

$$D = \{(\mathbf{x}_i, y_i)\}_{i=1}^N, \quad \mathbf{x}_i \in \mathbb{R}^d, \quad y_i \in \{0, 1\} \quad (1)$$

where $\mathbf{x}_i$ is the feature vector of the i -th job and y_i is the binary label, with $y_i = 1$ denoting a successful job and $y_i = 0$ denoting a failed job. The task is modeled as a supervised binary classification problem.

The pipeline is composed of four major modules:

A. Classical Preprocessing Module

This module prepares the dataset for effective training and evaluation. Missing values in resource features are imputed

using mean substitution:

$$x_i^{(\text{feature})} \leftarrow \frac{1}{N} \sum_{j=1}^N x_j^{(\text{feature})} \quad (2)$$

Feature values are then normalized using Min–Max scaling to map all features into the $[0, 1]$ range:

$$x_{\text{scaled}} = \frac{x - x_{\min}}{x_{\max} - x_{\min}} \quad (3)$$

To mitigate class imbalance, stratified sampling is applied to produce subsets (e.g., 5,000–10,000 samples) with balanced distributions of successful and failed jobs. This ensures that both classical and quantum models are trained fairly on balanced data.

B. Quantum Feature Embedding and Classification Module

Each normalized feature vector $\mathbf{x}_i$ is mapped into a quantum Hilbert space using a parameterized quantum feature map $\phi(\mathbf{x})$. We employ the ZZFeatureMap [12], which applies parameterized rotations and entangling gates to encode classical data into quantum states:

$$|\phi(\mathbf{x})\rangle = U_{\phi(\mathbf{x})}|0\rangle^{\otimes n} \quad (4)$$

The similarity between two job instances is measured using a quantum kernel:

$$K_Q(\mathbf{x}_i, \mathbf{x}_j) = |\langle \phi(\mathbf{x}_i) | \phi(\mathbf{x}_j) \rangle|^2 \quad (5)$$

This quantum kernel is blended with a classical Radial Basis Function (RBF) kernel to form a hybrid kernel:

$$K_{\text{hyb}} = (1 - \alpha)K_{\text{RBF}} + \alpha K_Q \quad (6)$$

where $\alpha \in [0, 1]$ controls the relative contribution of classical and quantum components. The hybrid kernel is then used to train a Quantum Support Vector Machine (QSVM) classifier.

C. Classical Benchmark Module

As a robust classical baseline, a Random Forest (RF) classifier is employed. Given an ensemble of T decision trees, the prediction is obtained by majority voting:

$$\hat{y}_{\text{RF}} = \arg \max_{c \in \{0,1\}} \sum_{t=1}^T I(h_t(\mathbf{x}) = c) \quad (7)$$

where $h_t(\cdot)$ is the prediction of the t -th tree and $I(\cdot)$ is the indicator function. Hyperparameters such as tree depth and number of estimators are tuned via cross-validation.

D. Hybrid Integration and Prediction Module

To enable flexible integration, predictions from the classical and quantum classifiers can be combined using a weighted ensemble strategy [14]–[16]. Specifically, given posterior probabilities from the quantum branch $P_{\text{quantum}}(c|\mathbf{x}_i)$ and the classical branch $P_{\text{classical}}(c|\mathbf{x}_i)$, the final decision is obtained as:

$$\hat{y}_i = \arg \max_{c \in \{0,1\}} \left(\alpha \cdot P_{\text{quantum}}(c|\mathbf{x}_i) + (1 - \alpha) \cdot P_{\text{classical}}(c|\mathbf{x}_i) \right) \quad (8)$$

Here, α serves as an adjustable parameter controlling the contribution of each model. While α can be optimized via validation (e.g., grid search, Bayesian optimization), the experiments in this paper fix $\alpha = 0.5$ and analyze the resulting hybrid QSVM performance.

This modular system design allows the framework to effectively capture complex workload characteristics in cloud environments. Classical preprocessing ensures high-quality inputs, the quantum kernel enriches similarity representations, and the hybrid integration enables trade-offs between classical robustness and quantum-enhanced generalization.

III. PERFORMANCE EVALUATION AND CRITICAL DISCUSSION

In this section, we present a comparative analysis of our proposed hybrid classical-quantum model, compared to a strong classical baseline for job failure prediction in distributed cloud environments. The experiments are based on a subset of the Alibaba Cluster Trace v2018 [17], where the objective is to predict the binary outcome of job execution using scheduling and resource allocation features. The evaluation was conducted on 5,000 samples, with balanced stratified splits (80/20 train-test), and metrics include Accuracy, F1 Score, ROC AUC, Average Precision (AP), and fit-time training cost.

Figure 2 presents the performance of RF and the Hybrid QSVM ($\alpha = 0.5$). The RF baseline exhibits excellent predictive capabilities, with an accuracy of 0.98, F1 score of 0.98,

and near-perfect ROC AUC and AP (≈ 1.0). In contrast, the Hybrid QSVM has lower predictive scores (Accuracy=0.80, F1=0.81, ROC AUC=0.88, AP=0.84), yet it remains competitive considering the more constrained feature space and the inclusion of quantum feature mappings.

A crucial aspect of this comparison lies in training efficiency. The RF model requires ≈ 0.22 s for fitting, whereas the Hybrid QSVM fit cost is ≈ 0.54 s. Importantly, these times reflect fit-only costs under unbiased experimental conditions, excluding heavy kernel precomputation for the QSVM. Although the QSVM does not outperform RF in predictive metrics, the findings illustrate an alternative trade-off: by leveraging quantum feature maps blended with classical kernels ($\alpha = 0.5$), QSVM is able to capture nonlinear resource-failure interactions that are otherwise modeled primarily through ensemble depth in RF. This establishes the complementary role of quantum components in improving model diversity, even if accuracy does not surpass classical ensembles at this scale.




Fig. 2. Performance comparison between Random Forest and Hybrid QSVM ($\alpha = 0.5$) on job failure prediction. While RF achieves the strongest predictive metrics, the hybrid QSVM demonstrates competitive performance and highlights the trade-offs between accuracy and computational efficiency in hybrid quantum-classical learning.

These findings highlight that RF dominates in predictive performance and remains the most reliable choice under current conditions. However, the Hybrid QSVM illustrates potential scalability benefits for evolving workloads, particularly where feature interactions become more nonlinear or classical ensembles face diminishing returns. Moreover, the hybrid approach demonstrates how quantum and classical kernels can be combined within a modular pipeline, ensuring that incremental improvements in quantum hardware and solvers can directly translate into stronger real-world applicability.

IV. CONCLUSION

This work proposed a hybrid classical-quantum framework for job failure prediction in distributed cloud environments, evaluated on a subset of the Alibaba Cluster Trace v2018 dataset. Results on 5,000 balanced samples show that RF achieves superior predictive performance, while the hybrid QSVM ($\alpha = 0.5$) offers competitive results with potential advantages in capturing nonlinear dependencies through

quantum-enhanced kernels. Although current quantum limitations constrain standalone performance, the hybrid design demonstrates how quantum and classical models can complement each other, laying the groundwork for scalable, real-time failure prediction in dynamic cloud workloads. Future efforts will focus on larger dataset sweeps, amortized training cost analysis, and optimizing hybrid integration strategies as quantum hardware advances.

REFERENCES

- [1] M. Armbrust, A. Fox, R. Griffith, A. D. Joseph, R. Katz, A. Konwinski, G. Lee, D. Patterson, A. Rabkin, I. Stoica *et al.*, "A view of cloud computing," *Communications of the ACM*, vol. 53, no. 4, pp. 50–58, 2010.
- [2] Q. Zhang, L. Cheng, and R. Boutaba, "Cloud computing: state-of-the-art and research challenges," *Journal of Internet Services and Applications*, vol. 1, no. 1, pp. 7–18, 2010.
- [3] A. Ganapathi, Y. Chen, A. Fox, R. Katz, and D. Patterson, "Statistics-driven workload modeling for the cloud," in *2010 IEEE 26th International Conference on Data Engineering Workshops (ICDEW 2010)*. IEEE, 2010, pp. 87–92.
- [4] M. Jassas and Q. H. Mahmoud, "Failure analysis and characterization of scheduling jobs in google cluster trace," in *IECON 2018-44th Annual Conference of the IEEE Industrial Electronics Society*. IEEE, 2018, pp. 3102–3107.
- [5] A. Verma, L. Pedrosa, M. Korupolu, D. Oppenheimer, E. Tune, and J. Wilkes, "Large-scale cluster management at google with borg," in *Proceedings of the tenth european conference on computer systems*, 2015, pp. 1–17.
- [6] A. Gulati, A. Holler, M. Ji, G. Shanmuganathan, C. Waldspurger, and X. Zhu, "Vmware distributed resource management: Design, implementation, and lessons learned," *VMware Technical Journal*, vol. 1, no. 1, pp. 45–64, 2012.
- [7] J. Gao, H. Wang, and H. Shen, "Task failure prediction in cloud data centers using deep learning," *IEEE transactions on services computing*, vol. 15, no. 3, pp. 1411–1422, 2020.
- [8] J. Biamonte, P. Wittek, N. Pancotti, P. Rebentrost, N. Wiebe, and S. Lloyd, "Quantum machine learning," *Nature*, vol. 549, no. 7671, pp. 195–202, 2017.
- [9] M. Schuld, I. Sinayskiy, and F. Petruccione, "An introduction to quantum machine learning," *Contemporary Physics*, vol. 56, no. 2, pp. 172–185, 2015.
- [10] J. Preskill, "Quantum computing in the nisq era and beyond," *Quantum*, vol. 2, p. 79, 2018.
- [11] C. Blank, A. J. Da Silva, L. P. de Albuquerque, F. Petruccione, and D. K. Park, "Compact quantum kernel-based binary classifier," *Quantum Science and Technology*, vol. 7, no. 4, p. 045007, 2022.
- [12] V. Havlíček, A. D. Córcoles, K. Temme, A. W. Harrow, A. Kandala, J. M. Chow, and J. M. Gambetta, "Supervised learning with quantum-enhanced feature spaces," *Nature*, vol. 567, no. 7747, pp. 209–212, 2019.
- [13] M. Benedetti, E. Lloyd, S. Sack, and M. Fiorentini, "Parameterized quantum circuits as machine learning models," *Quantum science and technology*, vol. 4, no. 4, p. 043001, 2019.
- [14] M. J. Hall and M. Reginatto, "Hybrid quantum-classical ensembles," in *Ensembles on Configuration Space: Classical, Quantum, and Beyond*. Springer, 2016, pp. 161–190.
- [15] S. McFarthing, A. Pillay, I. Sinayskiy, and F. Petruccione, "Classical ensembles of single-qubit quantum variational circuits for classification," *Quantum Machine Intelligence*, vol. 6, no. 2, p. 81, 2024.
- [16] M. Schuld and F. Petruccione, "Quantum ensembles of quantum classifiers," *Scientific reports*, vol. 8, no. 1, p. 2772, 2018.
- [17] A. Group, "Alibaba cluster trace data 2018," <https://github.com/alibaba/clusterdata>, 2018, accessed: 2025-08-16.