International Journal of Recent Development in Engineering and Technology
Website: www.ijrdet.com (ISSN 2347-6435 (Online) Volume 15, Issue 08, August 2026)

A Reproducible Hybrid Quantum–Classical Anomaly-Detection Pipeline for High-Dimensional Tabular Data under Ideal and Noisy Conditions

Dr. Manusankar C
PG & Research Department of Computer Science, S S V College, Valayanchirangara

***Abstract*— This paper presents a reproducible hybrid quantum–classical protocol for binary anomaly detection in high-dimensional tabular data. The protocol combines leakage-aware classical preprocessing, qubit-compatible representation learning, a bandwidth-controlled ZZ quantum feature map, and a precomputed-kernel Support Vector Machine (SVM). The fidelity kernel is specified using the correct adjoint sequence $U_\varphi^\dagger(x_i; \gamma)U_\varphi(x_j; \gamma)$, where the state $U_\varphi(x_j; \gamma)|0\rangle^{\otimes n}$ is prepared before the adjoint is applied. Because finite-shot estimation and NISQ noise can render an empirical Gram matrix indefinite, the protocol includes symmetrization, eigenspectrum inspection, PSD projection, ridge stabilization, and unit-diagonal normalization before dual SVM optimization. The evaluation design explicitly separates full-dimensional classical baselines from qubit-constrained baselines, so that the information cost of dimensionality reduction is not attributed to quantum processing. It also treats quadratic kernel construction as an operational scope condition: quantum-kernel experiments use a reported stratified subset, while low-rank landmark approximations are defined as a separate scalability study. No unexecuted empirical results are claimed; the manuscript is a rigorous experimental specification for ideal and noisy simulation studies.**

***Keywords*— anomaly detection, hybrid quantum–classical computing, quantum kernels, NISQ noise, positive semi-definite repair, dimensionality reduction, scalability**

I. INTRODUCTION

Anomaly detection seeks observations that depart from an expected distribution. It is central to fraud monitoring, network security, industrial sensing, and health-related screening. High-dimensional tabular data makes this task difficult because variables have heterogeneous scales, correlated structure, nonlinear interactions, and highly imbalanced labels. A useful experimental methodology must therefore evaluate both detection quality and the representational compromises made before a model is applied.

Quantum kernel estimation maps a classical vector to a quantum state and uses state similarity as a kernel value. This can be combined with a conventional kernel classifier without training a variational circuit.

Nevertheless, practical quantum-kernel studies face three related difficulties: the number of qubits limits input dimensionality, finite shots and device noise perturb the kernel matrix, and complete kernel construction scales quadratically with the number of samples. These constraints must be explicit rather than hidden by a reduced benchmark.

This study defines a reproducible protocol rather than reporting unexecuted performance. Its contributions are: (i) a corrected fidelity-kernel circuit algebra with a tunable bandwidth parameter γ; (ii) an explicit PSD restoration procedure for noisy and shot-estimated Gram matrices; (iii) a dual-space benchmark that evaluates classical algorithms on full and reduced representations; and (iv) a bounded-sample and optional low-rank scalability protocol for $O(N^2)$ kernel evaluation.

II. RELATED WORK

A. Classical Anomaly Detection

Classical tabular anomaly detection includes unsupervised approaches such as Isolation Forest and One-Class SVM as well as supervised classifiers when labels are available. Random Forest and RBF-SVM are useful supervised references because they respectively model nonlinear decision rules through ensembles and through a conventional kernel geometry. Their performance on unreduced features is necessary for determining whether a qubit-limited representation preserves sufficient task information.

B. Quantum Feature Spaces and Kernel Concentration

Quantum feature maps embed x into $|\psi(x)\rangle = U_\varphi(x)|0\rangle^{\otimes n}$. The resulting fidelity kernel is theoretically PSD under exact evaluation. However, expressive entangling maps may exhibit kernel concentration: entries become nearly uniform as qubit number or depth grows, weakening discriminative geometry. Controlled depth and a tuned input scaling parameter are consequently included in the present protocol.

C. NISQ Kernel-Matrix Validity

Hardware noise and finite-shot estimation cause an empirical similarity matrix to depart from its exact PSD form. Because standard dual SVM solvers assume a PSD kernel, indefinite matrices require explicit treatment. The proposed procedure therefore repairs the matrix and records the eigenspectrum before and after repair as a reproducibility diagnostic.

III. Theoretical Formulation

A. Bandwidth-Controlled Quantum Feature Map

Let $x \in R^d$ be a reduced input with $d = n$ qubits. After bounded normalization, a ZZ feature map with repetition depth $r \in \{1, 2\}$ and bandwidth $\gamma > 0$ is used:

$$U\varphi(x; \gamma) = [\ \prod_k RZ(2\gamma x_k)\ \prod(j,k)\in E\ RZZ(2\gamma(\pi-x_j)(\pi-x_k))\ H^{\otimes n}\ ]^r$$

Here $H^{\otimes n}$ forms a superposition, E denotes the selected connectivity graph, RZ is a single-qubit Z rotation, and RZZ is a two-qubit ZZ interaction. γ is selected inside training folds from a predefined grid; this prevents test-set information from influencing circuit selection and helps avoid uncontrolled phase scaling.

B. Fidelity Kernel and Correct Circuit Sequence

For xi and xj, define $|\psi(xi)\rangle = U\varphi(xi; \gamma)|0\rangle^{\otimes n}$. The quantum kernel is:

$$K(xi, xj) = |\langle\psi(xi)|\psi(xj)\rangle|^2 = |\langle 0|^{\otimes n}\ U\varphi^\dagger(xi; \gamma)\ U\varphi(xj; \gamma)\ |0\rangle^{\otimes n}|^2$$

The implementation first prepares $U\varphi(xj; \gamma)|0\rangle^{\otimes n}$ and then applies $U\varphi^\dagger(xi; \gamma)$. The probability of the all-zero outcome is the estimated fidelity. Reversing this operational description would apply the adjoint to the initial state rather than to the state prepared for comparison.

C. PSD Restoration for Empirical Gram Matrices

Let Kraw denote the finite-shot or noisy Gram matrix. The following deterministic procedure is applied only to the training Gram matrix, and the same fitted convention is used for test-kernel handling: (1) symmetrize, $Ksym = (Kraw + Kraw^T)/2$; (2) calculate $Ksym = V\Lambda V^T$; (3) set $\Lambda + = \mathrm{diag}(\max(0, \lambda 1), \ldots, \max(0, \lambda N))$; (4) reconstruct $KPSD = V\Lambda + V^T + \varepsilon IN$, with $\varepsilon = 10^{-5}$; and (5) normalize $Kij \leftarrow Kij/\mathrm{sqrt}(KiiKjj)$.

The count and magnitude of negative eigenvalues are recorded before repair.

D. Precomputed-Kernel SVM

For binary labels $yi \in \{-1,+1\}$, the repaired training kernel is supplied to a C-SVM with the dual objective:

$$\max\alpha\ \ \Sigma_{i=1}^{N}\ \alpha_i - \tfrac{1}{2}\ \Sigma_{i=1}^{N}\ \Sigma_{j=1}^{N}\ \alpha_i\alpha_j y_i y_j KPSD(xi,xj)$$

Subject to $0 \le \alpha i \le C$ and $\Sigma i\ \alpha iyi = 0$. C is tuned within the training data only. This formulation uses both label indices consistently and avoids applying an unverified noisy matrix directly to the solver.

IV. Proposed Hybrid Pipeline

A. Leakage-Aware Data Preparation

Each stratified training fold independently fits median imputation, robust interquartile-range scaling, feature selection, dimensionality reduction, circuit bandwidth, and SVM hyperparameters. The fitted transformations are then applied to the corresponding validation or test partition. Inputs sent to the circuit are bounded using $xnorm = \pi\ \tanh(xscaled)$, reducing extreme rotations.

B. Dual Representation Strategy

The design maintains two representations. The full representation dorig is retained for unconstrained classical baselines. A separate qubit-compatible representation $d = n \le 8$ is used for the quantum model and matched reduced-space baselines. Rather than treating PCA as anomaly preserving, the protocol regards compression as a possible source of task-relevant information loss. PCA may be combined with supervised feature ranking within the training fold, but neither technique is assumed to retain every low-variance or high-order interaction.

C. Ideal and Noisy Kernel Evaluation

The quantum kernel is first evaluated with an ideal statevector simulator. A noisy setting may then include specified gate, relaxation, and readout channels with a fixed shot budget. Every run reports the backend, noise parameters, shot count, qubit number, topology, circuit depth, γ, random seed, and wall-clock kernel-generation time.

International Journal of Recent Development in Engineering and Technology
Website: www.ijrdet.com (ISSN 2347-6435 (Online) Volume 15, Issue 08, August 2026)

TABLE I
REPRODUCIBLE PIPELINE CONFIGURATION

Stage	Specification
Data split	Stratified folds; all fitted transforms restricted to the training fold
Representations	Full dorig for classical references; reduced $d = n \leq 8$ for QSVM and matched baselines
Reduction	Training-fold PCA plus optional supervised ranking; compression penalty reported explicitly
Feature map	ZZ map, linear connectivity, $r \in \{1,2\}$, γ chosen on a predefined training-only grid
Kernel circuit	Prepare $U\varphi(xj;\gamma)\|0\rangle \otimes n$, then apply $U\varphi\dagger(xi;\gamma)$, then measure all-zero probability
PSD repair	Symmetrization, eigenvalue clipping, εIN ridge, and unit-diagonal normalization
Solver	Precomputed-kernel C-SVM with training-only selection of C
Scalability	Reported stratified subset for complete Gram matrix; landmark approximation evaluated separately

V. Experimental Methodology

A. Benchmark Matrix

All supervised experiments use identical stratified cross-validation partitions for models compared within a dataset. Evaluation reports precision, recall, macro F1-score, ROC-AUC where appropriate, kernel alignment, runtime, and the mean and standard deviation across folds. No result is reported until the associated configuration has been executed.

TABLE II
EXPERIMENTAL MATRIX

ID	Feature space	Model / purpose
E1-Full	Full dorig	RBF-SVM, Random Forest, and Isolation Forest: unconstrained classical reference
E1-Reduced	Reduced n	Same classical models: estimates the compression penalty
E2-Ideal	Reduced n	Ideal-statevector quantum-kernel SVM: compares quantum geometry with E1-Reduced
E3-Noisy	Reduced n	Noisy/shot-based quantum-kernel SVM with PSD repair
E4-Scale	$n \in \{2,4,6,8\}$; $r \in \{1,2\}$	Tests kernel concentration and depth/qubit sensitivity
E5-Bandwidth	Training-only γ grid	Tests scaling sensitivity and identifies non-collapsed kernels
E6-Approximate	Reported N and landmarks m	Compares complete and landmark/Nyström kernel construction as a separate scalability study

B. Quadratic Kernel Construction as a Scope Condition

A complete $N \times N$ training Gram matrix requires $N(N-1)/2$ distinct off-diagonal pair evaluations when symmetry is used, plus diagonal handling. Thus its circuit-evaluation burden is $O(N^2)$; with M shots per pair, measurement cost also scales with $M N^2$. The primary quantum-kernel experiments will therefore use a predetermined stratified subset whose N is selected before execution according to documented backend and shot resources. The manuscript will report Ntrain, Ntest, the number of kernel pairs, M, batching strategy, and elapsed time. This bounded subset is an experimental scope condition, not a claim of unrestricted scalability.

C. Scalable Extension

For a larger-sample study, the protocol defines a separate landmark experiment. A set of $m \ll N$ representative training points is selected within the training fold; circuit evaluations estimate cross-kernel and landmark-kernel blocks, and a low-rank Nyström reconstruction is used. The approximate setting is never pooled with full-Gram results without being labeled separately, because it changes both computation and the induced kernel.

D. Interpreting High-Dimensional Compression

The quantum model is evaluated only after reducing dorig to $n \leq 8$. Consequently, its performance is interpreted as detection on a qubit-compatible representation, not as direct processing of the complete high-dimensional problem. The gap between E1-Full and E1-Reduced quantifies the observed compression penalty. The gap between E1-Reduced and E2-Ideal identifies any additional effect associated with the quantum-kernel representation on the same reduced inputs.

International Journal of Recent Development in Engineering and Technology
Website: www.ijrdet.com (ISSN 2347-6435 (Online) Volume 15, Issue 08, August 2026)

VI. Results Reporting Protocol

The final empirical report will include: (i) a configuration table for every fold and backend; (ii) metric tables with mean and standard deviation; (iii) full-versus-reduced classical comparisons; (iv) ideal-versus-noisy quantum comparisons; (v) eigenvalue diagnostics before and after PSD repair; (vi) distributions or summary statistics of off-diagonal kernel entries for concentration analysis; and (vii) kernel-pair counts and timing data. Failed runs, solver warnings, and excluded samples must be reported with their causes.

VII. Discussion And Limitations

The pipeline resolves important algebraic and solver-validity concerns, but it does not eliminate physical constraints. First, complete quantum-kernel construction has $O(N^2)$ pairwise cost and is therefore suitable only for explicitly bounded sample sizes under the stated resources. Landmark approximations can reduce the evaluation burden, but they introduce an additional approximation that must be evaluated separately.

Second, mapping dorig $\gg$ 10 features to $n \leq 8$ quantum inputs necessarily omits information. Supervised ranking, nonlinear preprocessing, and PCA can prioritize useful signals, but none guarantees preservation of subtle low-variance anomaly cues or fine-grained interactions. Full-dimensional classical references are consequently not optional: they identify whether performance changes arise from the quantum model or from the representation bottleneck. Finally, a favorable result on reduced data would establish only a conditional advantage within that reduced representation, not a general quantum advantage for the original dataset.

VIII. Conclusion And Future Work

This manuscript defines a reproducible and mathematically consistent hybrid quantum–classical anomaly-detection protocol. It corrects the fidelity-circuit ordering, restores PSD validity for noisy Gram matrices, introduces tunable bandwidth control, and separates full-dimensional classical performance from qubit-constrained comparisons. By reporting bounded sample sizes and compression effects explicitly, the study avoids overstating the reach of present-day quantum kernels. Future work may investigate learned dimensionality reduction, adaptive feature maps, low-rank kernel estimation, and hardware executions with calibration-aware error mitigation.

REFERENCES

[1] V. Chandola, A. Banerjee, and V. Kumar, “Anomaly detection: A survey,” ACM Computing Surveys, vol. 41, no. 3, pp. 1–58, 2009.

[2] F. T. Liu, K. M. Ting, and Z.-H. Zhou, “Isolation forest,” in Proc. IEEE International Conference on Data Mining, 2008, pp. 413–422.

[3] B. Schölkopf et al., “Estimating the support of a high-dimensional distribution,” Neural Computation, vol. 13, no. 7, pp. 1443–1471, 2001.

[4] V. Havlíček et al., “Supervised learning with quantum-enhanced feature spaces,” Nature, vol. 567, pp. 209–212, 2019.

[5] M. Schuld and N. Killoran, “Quantum machine learning in feature Hilbert spaces,” Physical Review Letters, vol. 122, 040504, 2019.

[6] S. Thanasilp, S. Wang, M. Cerezo, and Z. Holmes, “Exponential concentration and untrainability in quantum kernel methods,” Quantum, vol. 7, p. 1153, 2023.

[7] H.-Y. Huang et al., “Power of data in quantum machine learning,” Nature Communications, vol. 12, p. 2631, 2021.

[8] J. Preskill, “Quantum computing in the NISQ era and beyond,” Quantum, vol. 2, p. 79, 2018.

[9] C. K. I. Williams and M. Seeger, “Using the Nyström method to speed up kernel machines,” in Advances in Neural Information Processing Systems, 2001.

[10] T. Hubregtsen et al., “Training quantum embedding kernels on near-term hardware,” PRX Quantum, vol. 3, 030327, 2022.