

Quantum Spectral Anomaly Detection

Yewei Yuan, Michele Minervini, Mark M. Wilde, and Nana Liu

Abstract

A core task in quantum anomaly detection is to quantify how strongly a test quantum state deviates from normal states. Classical principal component analysis (PCA) evaluates this anomaly score relative to the subspace spanned by selected leading eigenvectors, but for quantum data, centering, eigenvector recovery, Gram-matrix construction, or QRAM-style loading can be more expensive than estimating the anomaly score itself. We propose Quantum Spectral Anomaly Detection (QSPADE), which computes PCA-like anomaly scores directly from the spectrum of the average normal state. Instead of hard PCA rank selection, QSPADE uses a smooth, temperature T -controlled spectral threshold, allowing near-threshold components to contribute partially and making the score continuous and less sensitive to noise or arbitrary cutoffs. QSPADE recovers the hard-projector PCA score when $T \rightarrow 0$. The resulting measurement-based detector can be calibrated with sample complexity independent of the data dimension. Simulations show that QSPADE behaves like kernel PCA on encoded classical data and detects changes across a transverse-field Ising transition without predefined order parameters. Consequently, QSPADE gives an efficient framework for both quantum-kernel anomaly detection and the monitoring of quantum-native systems.

1 Motivation and contribution

Anomaly detection is a one-class learning problem: from examples assumed to be normal, one assigns a new test object an anomaly score so that larger scores indicate stronger deviation from the normal source. Classical PCA gives a spectral way to construct such scores. After centering the training vectors, PCA keeps the dominant eigenspace of the covariance matrix; the usual residual statistic Q is the squared weight of a test point outside this principal subspace, while Hotelling's $\mathcal{T}^2$ statistic measures unusually large leverage inside the retained subspace [1–3].

For quantum data, the output is still only a scalar anomaly score, but the standard PCA pipeline asks for more information than the score itself. Quantum kernel PCA and quantum one-class methods can import nonlinear feature spaces, while they commonly rely on pairwise Gram matrices or QRAM-style data access assumptions [4–6]. Quantum PCA and block-encoding methods give spectral access to density operators or matrices [7, 8], yet full eigenvector recovery or hard spectral resolution may be unnecessary if the only target is to decide whether a test state lies in the learned normal support. Variational one-class models and quantum autoencoders avoid explicit diagonalization, but replace it by trainable circuits (may face barren plateaus) and their optimization costs [9–11].

QSPADE is designed around the scalar nature of the task, avoiding the extra reconstruction, matrix access, rank selection, and optimization cost. Its main point is not to reconstruct the principal components of quantum data, but to measure the acceptance probability of a test state under a learned quantum detector. Moreover, classical data can be centered before being embedded into feature states if the goal is to reproduce centered PCA monitoring. However, quantum-native states are already normalized positive operators; subtracting a mean state generally leaves the physical state space. QSPADE therefore uses the uncentered average state as the normal spectral object.

Contributions are threefold. First, we formulate the anomaly detector as an optimal measurement effect, accepting normal states with high probability. Second, we replace the discontinuous retained rank by a smooth retained-mass calibration, which is the quantity that can be estimated directly from dimension-independent measurements. Third, we give quantum implementations for both global calibration and single-target deployment, and show that QSPADE behaves like kernel PCA on encoded classical data and works on quantum-native data even when no clear spectral gap identifies a rank.

2 QSPADE detector

Let $\rho_1, \dots, \rho_N$ be normal training states. Their average state $C := \bar{\rho} = \frac{1}{N} \sum_{i=1}^N \rho_i$ plays the role of a quantum covariance object. A measurement effect is an operator M satisfying $0 \preceq M \preceq I$; it accepts an input state σ with Born probability $\text{Tr}(M\sigma)$. If $C = \sum_j \lambda_j |u_j\rangle\langle u_j|$, the hard rank- K detector $P_K = \sum_{j \leq K} |u_j\rangle\langle u_j|$ maximizes normal acceptance $\text{Tr}(MC)$ among all rank- K effects. Thus the top eigenspaces of C are the quantum analogue of the principal support.

The problem with P_K is that the rank K is a discontinuous choice. When eigenvalues near the cutoff are close, whole spectral components flip from normal to anomalous under a tiny change in rank. QSPADE replaces this hard choice by a soft spectral detector

$$M_{\mu,T}^f = f\left(\frac{C - \mu I}{T}\right),$$

where $f : \mathbb{R} \rightarrow (0,1)$ is a monotone response function, for example a logistic or Gaussian-threshold response. The scalar μ is the spectral threshold and $T > 0$ is a temperature, or resolution, parameter. Eigenvectors of C with $\lambda_j \gg \mu$ are accepted with weight near one, eigenvectors with $\lambda_j \ll \mu$ are rejected with weight near zero, and boundary modes are accepted fractionally. As $T \rightarrow 0$, the detector converges to the hard projector onto the eigenspaces above μ . The soft detector is not only a heuristic smoothing of PCA. It is the unique solution of the regularized support-learning problem

$$\min_{0 \preceq M \preceq I} \{\mu \text{Tr}(M) - \text{Tr}(MC) + T \text{Tr} \phi_f(M)\},$$

where ϕ_f is the convex generator associated with the response f . The term $-\text{Tr}(MC)$ rewards acceptance of normal data, $\mu \text{Tr}(M)$ penalizes detector size, and $T \text{Tr} \phi_f(M)$ prevents an all-or-nothing occupation pattern near the spectral cutoff.

For a desired retained mass $\alpha \in (0,1)$, QSPADE chooses the unique threshold μ_α satisfying $\text{Tr}(M_{\mu_\alpha,T}^f C) = \alpha$. The residual anomaly score of a test state σ is then

$$Q_\alpha^f(\sigma) = 1 - \text{Tr}(M_{\mu_\alpha,T}^f \sigma).$$

This is the default QSPADE score. It is the quantum analogue of the PCA residual Q , but with a retained-mass threshold rather than a hard retained rank. A Hotelling-type score can also be built by taking differences of nearby soft thresholds, thereby forming soft spectral windows; this is useful when one wants to detect abnormal leverage inside the retained support.

3 Quantum implementation

The detector above can be accessed without explicitly recovering eigenvectors. One implementation uses a qumode control: choose the response f and prepare a control-qumode with reflected position density $g_{f,T_1}(q) = T_1^{-1} f'(-q/T_1)$. After a controlled evolution $U_C = e^{-i\hat{p} \otimes C/T_2}$, measuring the position of the control and applying a classical threshold implements $M_{\mu,T_1 T_2}^f$. Because changing μ changes only the classical threshold, the same measurement records can be reused to sweep many retained-mass targets. The empirical tail function of the position outcomes estimates $\text{Tr}(M_{\mu,T}^f C)$, and the Dvoretzky–Kiefer–Wolfowitz inequality gives dimension-independent

$$O(\eta^{-2} \log \delta^{-1})$$

shots to learn the retained-mass curve to additive error η with failure probability δ [12, 13]. In the sample-access setting, this type of evolution $e^{itC/T}$ on qubits can be achieved by density matrix exponentiation [7, 14] using copies of C .




Figure 1: **Left:** Average normal state spectrum and soft occupations, showing hard rank instability. Black points: eigenvalues $\lambda_j(C)$ (log scale); dotted line: sampling resolution $1/N$ ($N = 600$). Vertical dashed line: hard rank $K(\alpha) = 7$ ($\alpha = 0.99$). Colored curves (right axis): soft occupations $f((\lambda_j - \mu_\alpha)/T)$ for various T . **Right:** Mean anomaly scores for a dominant normal (50%), rare normal (1.6%), and paramagnetic anomaly mode. Q_{raw} and hard $K = 4$ projectors score the rare normal mode above the anomaly, while $K = 6$ yields zero, hiding its rarity. QSPADE yields an intermediate score; resolving the tail via T keeps the rare normal mode below the anomaly. Thus, soft scoring distinguishes normal rarity from genuine anomalies.

For a single target α , a qubit-only route is also possible. If the response f is approximated by a bounded superposition of phases, then $\text{Tr}[f((C - \mu I)/T)\rho]$ can be estimated by randomized Hadamard tests using controlled evolutions $e^{itC/T}$. The shift by μI contributes only a known scalar phase on the control qubit, so calibrating a different threshold does not require a different data-register simulation. This route has sample complexity $O(A_f^2 \eta^{-2} \log \delta^{-1})$, where A_f is the sampling prefactor of the chosen phase representation [15].

4 Numerical Results

We test QSPADE in two settings corresponding to its intended uses. First, for classical data embedded as quantum feature states, QSPADE behaves as a quantum kernel anomaly detector: on blob, moon, and ring geometries, QSPADE obtains ROC-AUC scores 0.93, 0.88, 0.88, comparable to RBF kernel-PCA (0.93, 0.87, 0.82) and much more stable than linear PCA (0.89, 0.66, 0.57) on curved data.

Second, we use quantum-native transverse-field Ising data. The normal source is a low-energy mixture from the ordered phase at field $h_A = 0.4$, while anomalies come from the paramagnetic phase at $h_B = 1.2$. The average normal state has a graded spectrum, so a hard retained rank is ambiguous: several nearby retained-mass targets choose different ranks, and the cutoff reaches the sampling-resolution floor. As shown in Fig. 1, QSPADE instead assigns fractional occupation to boundary modes and separates rare but valid normal sectors from genuine support mismatch.

The distinction is important operationally. A density-weighted raw overlap score can mark a rare normal mode as more anomalous than a paramagnetic test state, simply because the rare mode has small probability inside C . A hard projector has the opposite problem: depending on the chosen rank, the same rare mode is either fully rejected or fully absorbed. In our simulation, at the resolution that resolves the spectral tail, QSPADE gives this rare normal sector a small but nonzero score while assigning a substantially larger score to the paramagnetic anomaly class. Thus the score carries information about rarity within the normal support without confusing rarity with leaving the support altogether.

References

- [1] Ian T. Jolliffe. *Principal Component Analysis*. Springer Series in Statistics. Springer, 2 edition, 2002. doi: 10.1007/b98835.
- [2] Harold Hotelling. The generalization of Student’s ratio. *The Annals of Mathematical Statistics*, 2(3):360–378, 1931. doi: 10.1214/aoms/1177732979.
- [3] J. Edward Jackson and Govind S. Mudholkar. Control procedures for residuals associated with principal component analysis. *Technometrics*, 21(3):341–349, 1979. doi: 10.1080/00401706.1979.10489779.
- [4] Nana Liu and Patrick Rebentrost. Quantum machine learning for quantum anomaly detection. *Physical Review A*, 97(4):042315, 2018. doi: 10.1103/PhysRevA.97.042315.
- [5] Maria Schuld and Nathan Killoran. Quantum machine learning in feature hilbert spaces. *Physical Review Letters*, 122(4):040504, 2019. doi: 10.1103/PhysRevLett.122.040504.
- [6] Vojtěch Havlíček, Antonio D. Córcoles, Kristan Temme, Aram W. Harrow, Abhinav Kandala, Jerry M. Chow, and Jay M. Gambetta. Supervised learning with quantum-enhanced feature spaces. *Nature*, 567(7747):209–212, 2019. doi: 10.1038/s41586-019-0980-2.
- [7] Seth Lloyd, Masoud Mohseni, and Patrick Rebentrost. Quantum principal component analysis. *Nature Physics*, 10(9):631–633, 2014. doi: 10.1038/nphys3029.
- [8] András Gilyén, Yuan Su, Guang Hao Low, and Nathan Wiebe. Quantum singular value transformation and beyond: Exponential improvements for quantum matrix arithmetics. In *Proceedings of the 51st Annual ACM SIGACT Symposium on Theory of Computing (STOC)*, pages 193–204, 2019. doi: 10.1145/3313276.3316366.
- [9] Gunhee Park, Joonsuk Huh, and Daniel K. Park. Variational quantum one-class classifier. *Machine Learning: Science and Technology*, 4(1):015006, 2023. doi: 10.1088/2632-2153/acafd5.
- [10] Korbinian Kottmann, Friederike Metz, Joana Fraxanet, and Niccolò Baldelli. Variational quantum anomaly detection: Unsupervised mapping of phase diagrams on a physical quantum computer. *Physical Review Research*, 3(4):043184, 2021. doi: 10.1103/PhysRevResearch.3.043184.
- [11] Hyeondo Oh and Daniel K. Park. Quantum support vector data description for anomaly detection. *Machine Learning: Science and Technology*, 5(3):035052, 2024. doi: 10.1088/2632-2153/ad6be8.
- [12] Aryeh Dvoretzky, Jack Kiefer, and Jacob Wolfowitz. Asymptotic minimax character of the sample distribution function and of the classical multinomial estimator. *The Annals of Mathematical Statistics*, 27(3):642–669, 1956. doi: 10.1214/aoms/1177728174.
- [13] Pascal Massart. The tight constant in the Dvoretzky–Kiefer–Wolfowitz inequality. *The Annals of Probability*, 18(3):1269–1283, 1990. doi: 10.1214/aop/1176990746.
- [14] Shelby Kimmel, Cedric Yen-Yu Lin, Guang Hao Low, Maris Ozols, and Theodore J. Yoder. Hamiltonian simulation with optimal sample complexity. *npj Quantum Information*, 3(1):13, March 2017. ISSN 2056-6387. doi: 10.1038/s41534-017-0013-7.
- [15] Alexander He, Nana Liu, and Mark M. Wilde. Fermi–Dirac machines as quantizations of neurons, 2026.