Discover Quantum Science

RESEARCH Open Access

Markov chain decision-making with quantum-assisted entropy kernels

Or Peretz[1] and Michal Koren[1*]

*Correspondence:
Michal Koren
michal.koren@shenkar.ac.il
[1]School of Industrial Engineering and Management, Shenkar—Engineering, Design, Art, Anne Frank 12, 5252317 Ramat Gan, Israel

Abstract

Sequential anomaly detection requires summary statistics that remain stable under nominal dynamics yet react sharply to distributional shifts. We introduce QEK-DTMC, a hybrid quantum–classical sequential detector in which quantum circuits are used solely to estimate overlap-based kernel similarities, while entropy estimation, density smoothing, and decision-making are performed classically via a Bayesian two-state DTMC filter. Quantum overlaps derived from a parameterized ZZ feature map are smoothed by a boundary-aware kernel density estimator, yielding a Shannon entropy statistic that feeds a Bayesian posterior update on a two-state DTMC. This design explicitly accounts for finite-shot (quantum) measurement noise and kernel concentration, aligning representation, estimation, and decision within a unified sequential rule. We establish finite-sample guarantees for kernel estimation and entropy evaluation and prove the correctness of the DTMC recursion under calibrated thresholds. Empirical results on controlled drift scenarios show strong performance: AUROC of 0.90, AUPRC of 0.28, false-alarm rate < 0.01, and mean detection delay of 14 steps. Robustness analyses confirm graceful degradation under noise, missingness, and covariate shift, and debouncing reduces spurious triggers without undermining sensitivity. These findings demonstrate that entropy-aware quantum kernels can be systematically integrated into sequential Bayesian detectors, offering a statistically principled and compatible pathway for real-time anomaly detection.

Keywords Quantum computing, Anomaly detection, Markov chain, Kernel density estimation

1 Introduction

Modern data systems increasingly require online monitoring under drift, noise, and latency constraints, so that decisions must be updated sequentially rather than in batches. Beyond supervised and unsupervised learning, sequential detection methods accumulate evidence over time and trigger alarms when change is likely, providing normative procedures and error control for online use [39, 42, 51, 54, 58]. Detector effectiveness hinges on summary statistics that remain stable under nominal dynamics yet respond sharply to departures. Information-theoretic functionals such as entropy and related divergences quantify uncertainty and dissimilarity without requiring fully

specified likelihoods, which is advantageous under model misspecification or partial observability [22, 30, 34, 52].

Quantum machine learning (QML) contributes hardware-efficient tools for estimating such summaries. Parameterized circuits encode data into quantum states, and overlaps between states can be estimated by interference routines such as the SWAP test, yielding kernel values that support classification and change detection [17, 19, 23, 48, 49]. Two practical challenges arise: kernel concentration can erode discrimination when encodings are overly expressive or misaligned, and finite shot budgets introduce measurement noise that must be accounted for in downstream decisions [26, 31, 32, 50, 55]. Recent quantum routines for estimating entropy and related distributional measures suggest a way to realign representation with uncertainty structure [27–29, 41, 43].

We present QEK-DTMC, a quantum-statistical detector that couples an entropy-aware quantum kernel with a discrete-time Markov chain backbone for online detection of anomaly and change. Quantum overlaps derived from a parameterized feature map feed a boundary-aware density estimate on the unit interval, from which a Shannon entropy statistic is computed and passed to a Bayesian posterior-odds update on a two-state DTMC that encodes regime persistence. The design explicitly models finite-shot uncertainty and kernel concentration effects, aligning the quantum representation with the sequential decision rule and enabling the use of calibrated thresholds under resource constraints. Empirically, we evaluate detection delay, false-alarm behavior, calibration, and receiver operating characteristic (ROC)/precision–recall trade-offs under stressors relevant to deployment and then compare results against classical baselines under matched error budgets.

The remainder of the paper is organized as follows. Section 2 contains scientific background and related work. Section 3 details the QEK-DTMC procedure and its correctness, including complexity considerations and finite-sample error control. Section 4 describes the empirical setup (data, procedure, and evaluation metrics). The results of experiments validating our method are presented in Sect. 5. Last, Sect. 6 discusses the main conclusions and suggestions for future directions.

2 Literature review

Quantum computing has matured into a practical toolkit for accelerating algorithms and uncovering structure in data, with direct implications for statistical inference and machine learning [6, 11, 59]. By exploiting superposition, interference, and controlled entanglement, quantum computing operates in Hilbert spaces whose effective dimensionality surpasses classical representations [5, 38, 45]. However, noise and decoherence necessitate representations and estimators that are both expressive and statistically stable, particularly for online, non-independent and non-identically distributed (non-i.i.d.) decision problems [42, 43, 54]. QML surveys consistently emphasize that the choice of data encoding and feature-map circuit largely determines the induced kernel geometry and the resulting statistical power in kernel-based learning. In particular, angle and amplitude encoding schemes govern how distances in the input space are mapped to Hilbert-space overlaps and how much overlap variability remains under finite-shot estimation, both of which are central to detection [9, 23, 38, 49].

Classical quickest-detection theory provides the normative baseline for online anomaly detection. Under i.i.d. observations, the sequential probability ratio test (SPRT)

is optimal for binary hypotheses, employing likelihood-ratio stopping to minimize expected sample size at prescribed error levels [42, 58]. For persistent monitoring, the cumulative sum (CUSUM) procedure implements repeated one-sided tests [39] and enjoys minimax optimality under Lorden's criterion [37], whereas Bayesian formulations trigger when posterior odds cross a threshold [51, 54]. These Bayesian recursions induce a belief state that evolves as a discrete-time Markov chain (DTMC) driven by incoming evidence, making DTMCs a natural substrate for model-based detectors and for analyzing stopping-time distributions [42, 54]. In queueing and service systems, Bayesian inference and optimal measurement design have been used to estimate or control stochastic dynamics, offering templates for our setting [2, 40]. However, although these detectors are likelihood-centric, their efficacy in high dimensions hinges on informative statistics.

Entropy and divergence functionals (Shannon entropy, Kullback–Leibler divergence (KL), and Jensen-Shannon divergence (JSD)) quantify uncertainty and dissimilarity and serve as natural indicators of distributional shift (Shannon 1948) [30, 34]. In streaming settings, thresholding these functionals can control false alarms when estimators admit concentration. Increases in entropy often precede failures of downstream classifiers. When parametric likelihoods are misspecified, nonparametric approaches (such as density-ratio estimation, kernel two-sample tests, and entropy surrogates) provide robust detectors [18, 20, 22, 27, 28, 52]. From a representation viewpoint, detector power hinges on statistics that remain stable under the null hypothesis yet respond sharply to departures, motivating entropy-aware kernels to amplify weak shifts while preserving decision stability.

Quantum kernel methods implement circuit-based feature maps $\phi(x)$ and estimate inner products $k(x, x') = langle\phi(x), \phi(x') rangle$ by interference [23, 48, 49]. In particular, the SWAP test is the canonical primitive for fidelity/overlap estimation, with circuit variants and analyses that clarify statistical efficiency [17, 19]. A central practical issue is kernel concentration: overly expressive or misaligned constructions can collapse variability and erode discrimination. Theory quantifies the trade-offs and shows that highly expressive circuits may yield exponentially vanishing kernel variance [26, 55], echoing barren-plateau phenomena [11, 36]. Data-encoding results likewise demonstrate that misaligned encoders can induce concentration or vanishing gradients, reinforcing the need to preserve task-relevant variance under finite shot budgets [31, 32, 50].

Against this backdrop, entropy kernels become attractive: recent hardware-efficient routines estimate information-theoretic functionals central to detection with constant, low-depth circuits and accuracy consistent with classical baselines [27–29, 41]. Hence, such plug-and-play estimators enable the use of kernel definitions built from quantum-estimated entropies, conditional entropies, and mutual information, comparing stream windows via overlaps of the corresponding quantum states [12, 46].

Most quantum anomaly-detection studies are static or batch (e.g., quantum autoencoders, quantum support-vector machines (SVMs)) and demonstrate that appropriate quantum features can separate rare patterns [3, 7, 47]. Recent evidence suggests that quantum kernels can outperform classical detectors on realistic data [4], with complementary case studies described across scientific domains [24, 44, 60].

While many QML anomaly-detection demonstrations follow static or batch evaluation protocols, the algorithmic landscape extends beyond autoencoder and QSVM-based approaches [24]. Graph-centric formulations enable direct anomaly scoring without

supervised labels by assigning outlier scores to vertices through continuous quantum-walk dynamics initiated from a uniform superposition, with convergence guarantees and NISQ-oriented depth-mitigation variants [33]. More broadly, recent surveys organize quantum anomaly detection across supervised, unsupervised, and reinforcement-learning paradigms and emphasize that the "batch vs. streaming" distinction is largely orthogonal to the choice of model family [15, 57].

In parallel, a distinct body of work investigates sequential (early-stopping) quantum decision protocols, in which evidence is accumulated online and experiments terminate once a prescribed error budget is met [21]. This includes formal treatments of quantum sequential hypothesis testing that reduce the expected sample complexity relative to fixed-horizon tests, as well as extensions to continuously monitored quantum systems that require real-time inference from measurement streams [14]. These studies underscore that sequential quantum detection is an active research direction. By contrast, our focus lies on classical data streams, where quantum routines yield statistically controlled information-theoretic summaries (e.g., entropy-kernel overlaps that are integrated into a DTMC/Shiryaev-style Bayesian recursion for online anomaly and change detection under resource constraints [21, 35].

The sequential dimension remains underexplored: we need quantum summaries whose measurement noise admits finite-sample concentration compatible with online analysis, and a principled insertion of those summaries into Bayesian recursions in which the belief state evolves as a DTMC. This aligns directly with SPRT/CUSUM/Shiryaev-style updates by mapping quantum-derived statistics to likelihood-ratio or posterior-odds recursions [39, 42, 51, 54, 58]. In Markovian environments (such as queueing systems), such hybrids can surface weak entropy and divergence cues earlier while controlling false alarms [2].

Implementation leverages standard feature-map encodings (primarily angle encodings, with amplitude encodings used only when appropriate) because they are hardware efficient and induce overlap values in [0, 1] that can be estimated via a SWAP test. Encoding choice shapes both the geometry of the induced feature space and the finite-shot behavior of overlap estimation: for a fixed shot budget, the variance of the SWAP-based overlap estimator depends on the underlying overlap through the Bernoulli variance of the ancilla measurement outcome, so encodings that drive overlaps to concentrate near 0 or 1 can reduce discrimination or require more shots to reliably resolve small differences. Robust encoding design principles developed in supervised QML transfer naturally to the pre- versus post-change discrimination setting, motivating conservative, noise-tolerant feature maps in our sequential pipeline [31, 50]. The overarching design principle is measured expressivity: preserving task-relevant overlap variability while avoiding kernel concentration and feature collapse that can undermine statistical power under realistic shot budgets [26, 32].

3 Method

We introduce QEK-DTMC's quantum-statistical detector, a decoupled design in which a ZZ feature map with SWAP test yields pairwise similarities, boundary-aware kernel density estimation (KDE) produces a Shannon entropy signal, and a two-state DTMC performs a calibrated Bayesian belief update. The method is innovative in fusing an entropy-aware quantum kernel with a DTMC posterior, so that encoding, smoothing,

entropy computation, and stopping are co-designed and admit finite-sample guarantees. By explicitly controlling kernel concentration and shot noise and carrying that uncertainty into an SPRT-style posterior with tunable thresholds, QEK-DTMC delivers early, calibrated detection with hardware-efficient execution. Below, we first describe our notation and preliminary definitions. Second, we detail the quantum gates and operations for this method. Then, we proceed to describe the circuit architecture and implementation strategy and concludes with a formal correctness proof (Table 1).

3.1 Preliminaries

Our objective is to generate an alarm $a_t \in \{0, 1\}$, balancing early detection against false alarms, under a specified average run-length constraint. The four main components of the proposed algorithm are as follows:

1. *Stream and windowing.* Let $\{x_t\}_{t\geq 1} \subset mathbbR^d$ be an incoming stream. A data-driven rule, such that adaptive window (ADWIN), maintains a reference window $W_{t-1} = \{y_j\}_{j=1}^{W}$ of recent points prior to x_t. This window provides the classical context against which each incoming sample is compared, and it is updated online either through adaptive cuts or through a fixed-size first-in, first-out (FIFO) scheme, ensuring that the reference distribution reflects the most recent regime.
2. *Feature map and kernel.* A parameterized ZZ feature map $U_\Phi : mathbbR^d \to \mathrm{U}\,(2^m)$ encodes $x \mapsto |\phi\,(x)\,\rangle = U_\Phi\,(x)\,|0\rangle^{\otimes m}$. For $x, y \in mathbbR^d$, define the overlap kernel:

$$k\,(x, y) = |langle\phi\,(x)\,midphi\,(y)\,rangle|^2 \in [0, 1] \tag{1}$$

At time t we form the multiset $\{k\,(x_t, y_j) : y_j \in W_{t-1}\}$, which captures how the incoming state overlaps with recent history; this set becomes a sufficient summary for subsequent dispersion estimation.

Table 1 Notation used in this study

Symbol	Remarks
X_t	Sequential data stream at time t
W_t	Sliding reference window at time t
$\|xrangle$	Quantum-state encoding of data vector x
Φ_{ZZ}	Parameterized ZZ feature map
$k\,(x_i, x_j)$	Overlap kernel via SWAP test between states $\|x_irangle,\| x_jrangle$
S_t	Multiset of similarities at time t
h	Bandwidth parameter for beta-kernel KDE
$\hat{f}\,(s)$	Boundary-aware beta-kernel density estimate
H_t	Shannon entropy statistic at time t
π	Gauss–Legendre quadrature order for entropy evaluation
M	Quadrature nodes used in entropy approximation
L_t	Log-likelihood ratio at time t
$Z_t \in \{0, 1\}$	Latent DTMC state: 0 = no drift, 1 = drift
P	Transition probability matrix of two-state DTMC
γ_t	Posterior drift probability $\Pr\,(Z_t = 1midH_{1:t})$
α, β	Error-control parameters
τ	Decision threshold on posterior probability
d	Window size
N	Number of SWAP-test shots per pair
$\hat{k}_{ij}$	Empirical similarity estimator from SWAP tests
r	Debounce run length (consecutive time steps with $\pi_t \geq \tau$ required to raise an alarm)

3. *Boundary-aware KDE and entropy*. Since $k \in [0,1]$, we adopt a beta-kernel KDE on the unit interval,

$$\hat{f}_t(u) = \frac{1}{W}\sum_{j=1}^{W} K_h(u; s_j), \qquad s_j := \hat{k}(x_t, y_j) \tag{2}$$

where each $K_h(\cdot; s)$ is a beta density with mean s and dispersion parameter $h \in (0,1)$. The entropy statistic is

$$\hat{H}_t = -\int_0^1 \hat{f}_t(u)\, \log \hat{f}_t(u)\, du \tag{3}$$

Evaluated numerically by Gauss–Legendre quadrature on $[0,1]$ with order L [53]. We refer to this estimator as boundary aware because beta kernels have support exactly on $[0,1]$, ensuring that the density estimate does not leak probability mass outside the feasible similarity range and does not require ad hoc reflection or truncation corrections near zero and one [13]. This property is particularly important here because SWAP-test similarities are intrinsically bounded. This construction respects the bounded support of similarities and yields a smooth, strictly positive density, so the differential entropy is well defined and numerically stable under standard quadrature rules.

We adopt a boundary-aware beta-KDE on $[0,1]$ with an analytic, window-tied bandwidth schedule, $h(W) = \mathrm{clip}_{[h_{\min}, h_{\max}]}\left(\frac{c_h}{\sqrt{W}}\right)$, for $c_h > 0$. This schedule tracks the classical $W^{-1/2}$ bias-variance frontier while preventing over-smoothing.

4. *DTMC detection backbone*. We model H_t with a Gaussian family $p(H_t mid S_t = s) = mathcal N\left(\mu_s, \sigma_s^2\right)$ for $s \in \{0,1\}$. Offline calibration fits (μ_0, σ_0) on a pre-drift slice and (μ_1, σ_1) on a post-drift slice. Let $S_t \in \{N, D\}$ denote the latent state (no-drift/drift):

$$\Pr(S_t = D S_{t-1} = N) = \alpha, \qquad \Pr(S_t = N S_{t-1} = D) = \beta \tag{4}$$

With self-transition probabilities $1-\alpha$ (stay N) and $1-\beta$ (stay D). Writing $\pi_t = \Pr\left(S_t = D mid \hat{H}_{1:t}\right)$, the one-step prior is

$$\tilde{\pi}_t = (1-\beta)\,\pi_{t-1} + \alpha\,(1-\pi_{t-1}) \tag{5}$$

Given likelihoods g_0, g_1 for $\hat{H}_t$ under N, D, the posterior is

$$\pi_t = \frac{g_1\left(\hat{H}_t\right)\tilde{\pi}_t}{g_1\left(\hat{H}_t\right)\tilde{\pi}_t + g_0\left(\hat{H}_t\right)(1-\tilde{\pi}_t)} \tag{6}$$

We calibrate g_0, g_1 either offline from held-out pre/post-change segments or online via exponentially weighted parametric fits. The posterior π_t then drives decisions by thresholding, and the DTMC transitions regularize the belief dynamics and implicitly encode the persistence of regimes.

3.2 Quantum gates

We employ a minimal set of standard gates as the elementary building blocks of the QEK-DTMC detector. Each data vector x_t is encoded into an n-qubit quantum state via a depth-d ZZ feature map:

$$U_{\mathrm{rot}}^{(t)}(x) = \prod_{j=1}^{n} R_Z\left(\phi_j^{(t)}(x)\right)$$

$$U_{\mathrm{ent}}^{(t)}(x) = \prod_{(j,k)\in mathcal E} \exp\left(\frac{i}{2}\,\theta_{jk}^{(t)}(x)\; Z_j Z_k\right)$$

$$U_{\Phi}(x) = \left(\prod_{t=1}^{d_{\mathrm{ZZ}}} U_{\mathrm{ent}}^{(t)}(x)\; U_{\mathrm{rot}}^{(t)}(x)\right) \cdot \left(\otimes_{j=1}^{n} H_j\right)$$

where n is the number of qubits, d_{ZZ} is the number of ZZ blocks, and E denotes the chosen entangling-edge set (e.g., linear connectivity). Kernel similarities between encoded states are obtained using the standard SWAP test (Fig. 1), a canonical primitive for estimating state overlap (fidelity) via an ancilla measurement. Let $|\psi$ and $|\phi$ be two states. The procedure prepares $|0$ on an ancilla a, applies a Hadamard H to a, performs a controlled SWAP on registers ψ and ϕ controlled by a, and concludes with a second Hadamard on a before measurement. The probability of measuring 0 on the ancilla satisfies

$$p(0) = \frac{1}{2}\left(1 + |\psi|\phi|^2\right)$$

Hence the unbiased similarity estimator is $\hat{k} = 2\hat{p}(0) - 1$, where $\hat{p}(0)$ is the empirical mean outcome over N shots. Hoeffding's inequality provides the finite-sample guarantee

$$\Pr\left(\left|\hat{k} - k\right| \geq \varepsilon\right) \leq 2\exp\left(-\frac{N\varepsilon^2}{2}\right)$$

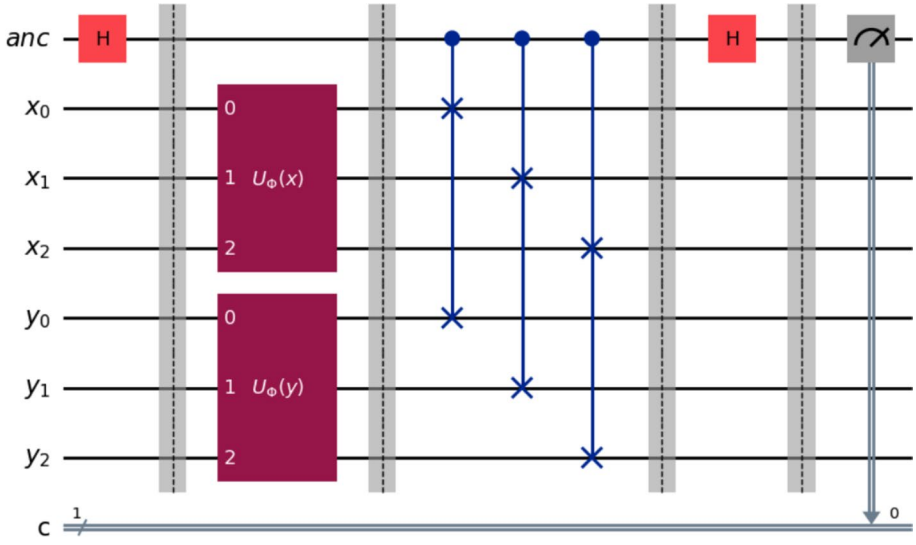

Fig. 1 Quantum SWAP-test circuit used to evaluate pairwise kernel similarities

3.3 Procedure flow and logic

The method is organized around four modular subroutines as follows. Each is well-defined, is self-contained, and provides explicit guarantees for statistical or computational accuracy.

1. *SwapTestKernel*. This routine prepares the quantum states $U_\Phi(x_t)$ and $U_\Phi(x_j)$, attaches an ancilla qubit, and executes a sequence of SWAP tests. Shots are batched, and execution halts once Hoeffding's bound ensures tolerance ε, $\Pr\left(\left|\hat{k}-k\right| \geq \varepsilon\right) \leq 2\exp\left(-2N\varepsilon^2\right)$, or once the shot cap $N_{\max}$ is reached. The subroutine returns the similarity estimate $\hat{k}(x_t, x_j)$ and the realized number of shots. This ensures concentration while avoiding oversampling.
2. *BetaKDEEntropy*. This routine transforms the similarity multiset K_t into a smooth density over $[0,1]$ using beta kernels:

$$\hat{f}(u) = \frac{1}{|K_t|} \sum_{k \in K_t} \mathrm{beta}\left(u \mid \alpha(k,\tau),\, \beta(k,\tau)\right)$$

where we parameterize the beta kernels by a single dispersion $\tau > 0$ via $\alpha(k,\tau) = \tau\, k$ and $\beta(k,\tau) = \tau\ (1-k)$. The Shannon entropy is then $\hat{H}_t = -\int_0^1 \hat{f}(u) \log \hat{f}(u)\ du$, evaluated numerically using Gauss–Legendre quadrature of order L.

3. *DTMCPosterior*. This routine maintains the Bayesian belief state. Given a prior π_{t-1} and likelihoods $L_0\left(\hat{H}_t\right)$ and $L_1\left(\hat{H}_t\right)$, the posterior recursion is

$$\widehat{L_1} = L_1\left(\hat{H}_t\right) \cdot \left((1-\beta)\,\pi_{t-1} + (1-\alpha)\,(1-\pi_{t-1})\right)$$

$$\pi_t = \frac{\widehat{L_1}}{L_0\left(\hat{H}_t\right) \cdot \left(\beta\,(1-\pi_{t-1}) + \alpha\pi_{t-1}\right) + \widehat{L_1}}$$

For numerical robustness, the recursion is carried out in log-odds form:

$$l_t = \log\left(\frac{\pi_t}{1-\pi_t}\right)$$

4. *WindowUpdate*. This routine maintains the context set W_t. Under ADWIN, cuts are applied adaptively to retain the most recent homogeneous segment. Under a fixed FIFO policy, the oldest sample is discarded as x_t is appended. In either case, W_t provides the reference distribution against which new observations are compared.

At initialization, a bootstrap segment seeds the window W_0, sets the prior π_0, precomputes quadrature nodes and weights, and calibrates likelihood models if desired. Thereafter, the system advances online in perpetuity. For window maintenance, WindowUpdate uses either a FIFO of fixed size W or an adaptive ADWIN-style change detector with significance level δ_{win} controlling the false-cut rate; we report which option is used in each experiment.

At each step, the new sample x_t first updates the reference window via WindowUpdate. Quantum similarities against all $x_j \in W_{t-1}$ are then estimated through SwapTestKernel, yielding the multiset K_t. This multiset is passed to BetaKDEEntropy, which produces a smoothed density and the entropy statistic $\hat{H}_t$. Finally, DTMCPosterior integrates $\hat{H}_t$ with the prior belief, producing the posterior π_t and updating the log odds. An alarm decision is triggered when π_t exceeds a pre-specified threshold τ. Optionally, we apply debouncing: an alarm is issued only if $\pi_t \geq \tau$ holds for r consecutive time steps, thereby suppressing isolated spikes (Fig. 2).

3.3.1 Notes

- We pick $h = \frac{c_h}{\sqrt{W}}$, with c_h chosen on a small grid during calibration and clipped to $[0.02, 0.20]$. Densities are clipped below by $\varepsilon_f \in \left[10^{-9}, 10^{-6}\right]$.
- Unless otherwise stated, we use $(c_h, h_{\min}, h_{\max}) = (0.4,\ 0.02,\ 0.20)$, which yields $h \in [0.02, 0.05]$ for $W \in [64, 256]$. For stability near the boundaries, densities are clipped below by ε_f as in Sect. 3.5.

SwapTestKernel$(x_t, W, N_0, \varepsilon, \delta)$

- $K_t \leftarrow \emptyset$
- **For each** $y \in W$
  - $N \leftarrow 0, \bar{Z} \leftarrow 0$
  - **While** $N < N_0$
    - SWAP(x_t, y)
    - update $\bar{Z}$, N
    - **If** $2\exp(-\frac{N\varepsilon^2}{2}) \leq \frac{\delta}{|W|}$ **Then** break
  - $\hat{k} \leftarrow 2\bar{Z} - 1$
  - $K_t \leftarrow K_t \cup \{\hat{k}.$
- **Return** $K_t, \max_{y \in W} N$

QEK-DTMC(W_0, π_0)

- $W \leftarrow W_0$, $\pi \leftarrow \pi_0$, $c \leftarrow 0$
- For each time-step t
  - $W \leftarrow W \cup \{x_t\}$ by **WindowUpdate**
  - $K_t, N_t \leftarrow$ **SwapTestKernel**$(x_t, W, N_0, \varepsilon, \delta)$
  - For $k_j \in K_t$
    - $s_j \leftarrow \frac{1+\hat{k}_j}{2}$
    - Define $H_t, \hat{f}_t(u)$
  - $\pi \leftarrow$**DTMCPosterior**$(\pi_{t-1}, \alpha, \beta, L, U)$
  - **If** $\pi \geq \tau$ **Then** $c \leftarrow c + 1$ **Otherwise** $c \leftarrow 0$
  - **Return** $c \geq \tau$

DTMCPosterior$(\pi_{t-1}, \alpha, \beta, L, U)$

- $\tilde{\pi} \leftarrow (1-\beta)\pi_{t-1} + \alpha(1-\pi_{t-1})$
- Evaluate likelihoods
- Update posterior π_t
- **If** $\pi_t \leq L$ **Then** use $H_t \mid S_t = 0$ to update
- **If** $\pi_t \geq U$ **Then** use $H_t \mid S_t = 1$ to update
- **Return** π_t

3.4 Cost and computational complexity

Each time step involves both quantum and classical components. On the quantum side, the gate depth scales as $O\left(W\ (d_\Phi + m)\right)$, where W is the window size, d_Φ the

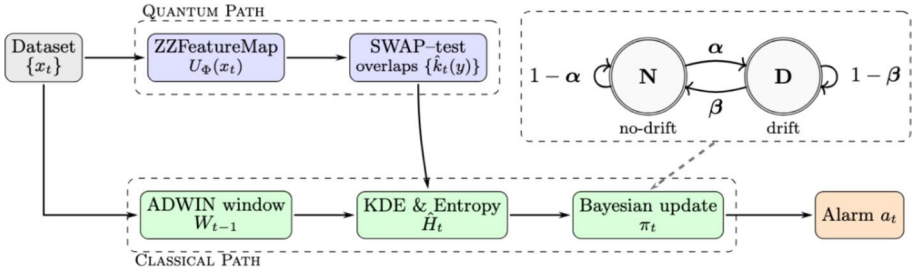

Fig. 2 Overall architecture of the QEK-DTMC detector

feature-map depth, and m the controlled-SWAP depth. The number of measurements scales as $O(WN)$, with realized shots $N \leq N_{\max}$. On the classical side, entropy evaluation via quadrature requires $O(WL)$ floating-point operations. To control overhead, shot counts are batched geometrically ($N_b \in \{16, 32, 64, \ldots\}$), preventing overshooting of the Hoeffding bound. Kernel estimates are clipped to a safe lower bound before applying the logarithm. Alarm signals are debounced using a run-length rule, requiring $\pi_t \geq \tau$ for r consecutive steps. These operational safeguards stabilize numerical behavior and suppress spurious transients without altering the statistical semantics of detection.

Resource scaling remains linear in the effective context size. Explicit encodings and pairwise SWAP tests against a classical window yield a per-step cost of $O(W\ (d_\Phi + m))$ gates and $O(WN)$ measurements. Memory usage is $O(W)$ for both the reference window and the similarity multiset.

3.5 Correctness

Let $\{\hat{k}_j\}_{j=1}^{W}$ be the SWAP-test overlap estimates in the reference window, each computed with N i.i.d shots. For a fixed pair (x_t, y), the SWAP ancilla outcome is Bernoulli distributed with $p_0 = \frac{1}{2}(1 + k(x_t, y))$ and estimator $\hat{k} = 2\hat{p}_0 - 1$, where $\hat{p}_0$ is the sample mean [16]. By Hoeffding's inequality for bounded i.i.d. variables [25], it follows that:

$$\Pr\left(|\hat{k} - k\rangle \varepsilon\right) \leq 2\exp\left(-\frac{N\varepsilon^2}{2}\right)$$

A union bound across the W comparisons imply that, for $N \geq \frac{2}{\varepsilon^2}\log\frac{2W}{\delta}$ we have $\max_{1\leq j\leq W}\left|\hat{k}_j - k_j\right| \leq \varepsilon$ with probability at least $1-\delta$ [8].

Let $\hat{f}$ denote the boundary-aware beta-kernel density on $[0,1]$ constructed from $\{\hat{k}_j\}_{j=1}^{W}$ [13] and clipped below at $\varepsilon_f \in (0,1)$. Let $\hat{H}$ be the corresponding entropy, evaluated by a quadrature of order L with error $E_{\mathrm{q}}(L)$ [56]. Assume that the kernel $K_h(u;s)$ is Lipschitz in its center s uniformly in $u \in [0,1]$ with constant L_h. Then, if $f^\star$ and $H_t^\star$ denote the density and entropy built from the exact $\{k_j\}$,

$$\int_0^1 |\hat{f}(u) - f^\star(u)|du \leq L_h\varepsilon$$

and, since $\varphi(x) = -x\log x$ is C_φ-Lipschitz on $[\varepsilon_f, 1]$ with $C_\varphi = 1 - \log\varepsilon_f$ (),

$$\left|\hat{H} - H_t^\star\right| \leq C_\varphi L_h\varepsilon + E_{\mathrm{q}}(L)$$

Hence finite-shot noise at the overlap level induces a directly tunable error at the statistic used by the filter. Let $\tilde{\pi}_t$ be the DTMC prior prediction from transition matrix $\begin{pmatrix} 1-\alpha & \alpha \\ \beta & 1-\beta \end{pmatrix}$, and let π_t be the posterior after observing $\hat{H}_t$, with emission densities g_0, g_1 for states $\{N, D\}$. Writing $l(p) = \log\left(\frac{p}{1-p}\right)$, the binary-state Bayesian filter updates additively in log odds [10]:

$$l(\pi_t) = l(\tilde{\pi}_t) + \log\left(\frac{g_1\left(\hat{H}_t\right)}{g_0\left(\hat{H}_t\right)}\right)$$

If the log-likelihood ratio $\log\left(\frac{g_1}{g_0}\right)$ is C_g-Lipschitz over the operating range, then with $H_t^\star$ the entropy from exact similarities and $\pi_t^\star$ the corresponding posterior,

$$\left| l\left(\pi_t\right) - l\left(\pi_t^\star\right) \right| \leq C_g \left| \hat{H}_t - H_t^\star \right| \leq C_g \left(C_\varphi L_h \, \varepsilon + E_{\mathrm{q}}\left(L\right) \right)$$

Because $l'\left(p\right) = \left[p\left(1-p\right)\right]^{-1}$, perturbations in log odds translate locally to probability scale with factor $\left[\,\tau\left(1-\tau\right)\,\right]^{-1}$ near a decision threshold $\tau \in (0,1)$. Consequently, whenever $|\pi_t^\star - \tau\,| \geq \gamma \rangle\, 0$ and

$$\left| \hat{H}_t - H_t^\star \right| \leq \frac{\gamma}{C_g} \min\left\{ 1, \frac{1}{\tau\left(1-\tau\right)} \right\}$$

The induced decision remains unchanged. To attenuate transient spikes, a run-length requirement of r consecutive posteriors above τ reduces the false-alarm probability by approximately the factor $\Pr(\pi_t \geq \tau)^r$ under near-independent short-run fluctuations, while adding at most $r-1$ steps of delay once drift persists.

4 Empirical study

We undertook a controlled empirical study to assess the proposed QEK-DTMC under synthetic data streams designed to exhibit drifts. The study emphasizes auditable procedure, explicit factor control, and reproducibility, with results reported as means and uncertainty intervals and supported by formal significance testing where applicable. The factorial design crosses four dimensions:

1. Drift mechanism (abrupt, gradual with widths $w \in \{50, 150\}$, and recurring).
2. Ambient dimension $d \in \{2, 8\}$.
3. Class imbalance $\{50/50,\ 30/70\}$.
4. Signal-to-noise ratio $(\mathrm{SNR}) \in \{5, 10, 20\}$.

For each grid point, we instantiate labeled timelines that mark nominal segments and drift neighborhoods, enabling matched evaluation across methods and posterior calibration of π_t. Streams are generated with deterministic seeds (20 paired seeds per setting), ensuring paired comparisons and variance reduction in statistical testing.

4.1 Experimental procedure

1. We vary the number of ZZ feature-map blocks $d_{ZZ} \in \{2, 4, 6, 8\}$, the feature-map depth proxy $d_\Phi \in \{1, 2, 3\}$, and the SWAP-test shot budget $N \in \{32, 64, 128, 256, 512\}$.
2. The context window is $W \in \{64, 128, 256\}$; the overlap-domain beta-KDE bandwidth is $h \in \{0.02, 0.05, 0.1\}$; and the entropy quadrature uses Gauss–Legendre order $L \in \{32, 64, 128\}$.
3. The prior parameters sweep is $(\alpha, \beta) \in \{0.01, 0.05, 0.1\}^2$; the decision threshold is $\tau \in \{0.6, 0.7, 0.8\}$; and the debounce run length is $r \in \{1, 3, 5\}$.
4. To probe stability, we introduce:

   (a) Quantum-noise aggregated to total error probabilities $\{0, 0.5\%\ , 1\%\ , 2\%\ \}$.
   (b) Observation missing $\{0\%\ , 5\%\ , 10\%\ , 20\%\ \}$.
   (c) Covariate shift (train on regime A, test on shifted regime B). Each stress setting reuses the base seeds to isolate treatment effects.

5. Each configuration yields a time series of entropy statistics H_t, anomaly posteriors π_t, and binary alarms. We summarize operating characteristics by detection-delay distributions (and their empirical cumulative distribution functions), mean detection delay (MDD), false-alarm rate (FAR) as a function of τ, and calibration quality via expected calibration error (ECE) and Brier score.

To characterize estimator quality, we also report overlap-estimation error $\left|\hat{k} - k\right|$ as a function of N (with concentration bounds where applicable), bandwidth-induced bias-variance trade-offs for h, and quadrature accuracy versus order L. For readers who want a fully worked numerical example (one and two qubit sanity-check of the pipeline), we provide an illustrative walkthrough in "Appendix".

To assess external validity beyond controlled synthetic drifts, we evaluate QEK-DTMC on a real streaming anomaly dataset from the Numenta Anomaly Benchmark [1]. We use the "temperature_system_failure" stream (real known cause), a univariate time series with labeled anomalous intervals. The stream is processed sequentially: at each time step t, a feature vector x_t is formed from a sliding window of recent observations and passed through the same pipeline used in the synthetic study. The signal is windowed with length $W = 128$ and standardized using statistics from an initial nominal segment; missing values are handled by forward fill followed by median imputation. We run QEK-DTMC with SWAP-shot budget $N = 128$, feature-map depth proxy $d_\Phi = 1$, beta-KDE bandwidth $h = 0.05$, quadrature order $L = 64$, fixed decision threshold τ, and debouncing run length $r = 3$.

5 Results

To evaluate the proposed QEK-DTMC detector, we conducted a series of experiments designed to assess both statistical accuracy and detection efficacy under controlled drift scenarios. First, we validate the concentration properties of the quantum kernel estimator. Second, we examine the stability and bias of the entropy statistic across synthetic benchmarks. Finally, we assess the detection power and false-alarm behavior of the DTMC-based posterior recursion. This sequence of experiments ensures that each algorithmic component is validated in isolation before the integrated procedure is analyzed.

5.1 Aggregate accuracy and timeliness

Core metrics (mean with dispersion) are summarized in Table 2. On average, the detector attains an area under the ROC curve (AUROC) = 0.902 (0.036) and area under the precision–recall curve (AUPRC) = 0.284 (0.072), with FAR = 0.004 (0.003) and MDD = 14.0 (6.2). Observed ranges indicate consistently strong ranking performance (AUROC 0.842–0.946) and operating-point-dependent precision–recall (AUPRC 0.190–0.445). False alarms are near zero for most settings (0–0.064), while delay remains bounded (4–28 steps). These patterns align with the expected latency–stability trade-off

Table 2 Core aggregate metrics of QEK-DTMC detector

Metric	Mean (dispersion)	Observed range
AUROC	0.902 (0.036)	[0.842, 0.946]
AUPRC	0.284 (0.072)	[0.190, 0.445]
FAR	0.004 (0.003)	[0, 0.064]
MDD	14.0 (6.2)	[4, 28]

governed by the debouncing parameter: tighter debouncing suppresses spurious triggers (lower FAR) at the cost of longer delay (higher MDD), whereas looser debouncing reduces delay but introduces more variability into the precision–recall relationship. Notably, the variability in timing (MDD) is larger relative to its mean than the variability in AUPRC, suggesting that debouncing primarily regularizes when alarms fire rather than how scores are calibrated.

5.2 Operating points and temporal behavior

We vary the debounce run length r to expose decision dynamics. Increasing r regularizes triggers in the expected direction: when $r = 1$, MDD = 16 and FAR = 0.04, whereas when $r = 3$, MDD = 56 and FAR = 0.021. Thus, numerically, raising r from 1 to 3 approximately halves false alarms while incurring additional delay. Time-aligned triggers are presented in Fig. 3, showing that debouncing suppresses transient excursions above threshold without altering steady-state decisions under persistent drift.

5.3 Calibration and reliability

Aggregate calibration indicates room for light post-hoc mapping: ECE = 0.158 and Brier = 0.211. Reliability curves (Fig. 4) show mild misalignment near the operating threshold, the regime in which per-stream Platt or isotonic remapping typically stabilizes FAR at fixed τ across covariate conditions. Hence, modest calibration suffices to align nominal and empirical trigger rates.

5.4 Robustness to noise and shift

Stress tests show robust behavior with predictable patterns. With shot-level noise increasing from 0 to 1%, performance declines slightly (AUROC 0.900 → 0.880; AUPRC 0.280 → 0.260), consistent with minor contamination reducing rank separation while leaving precision–recall largely intact. Under missingness, a 10% drop increases sparsity and yields a modest degradation (AUROC 0.904 → 0.865, AUPRC 0.280 → 0.251),

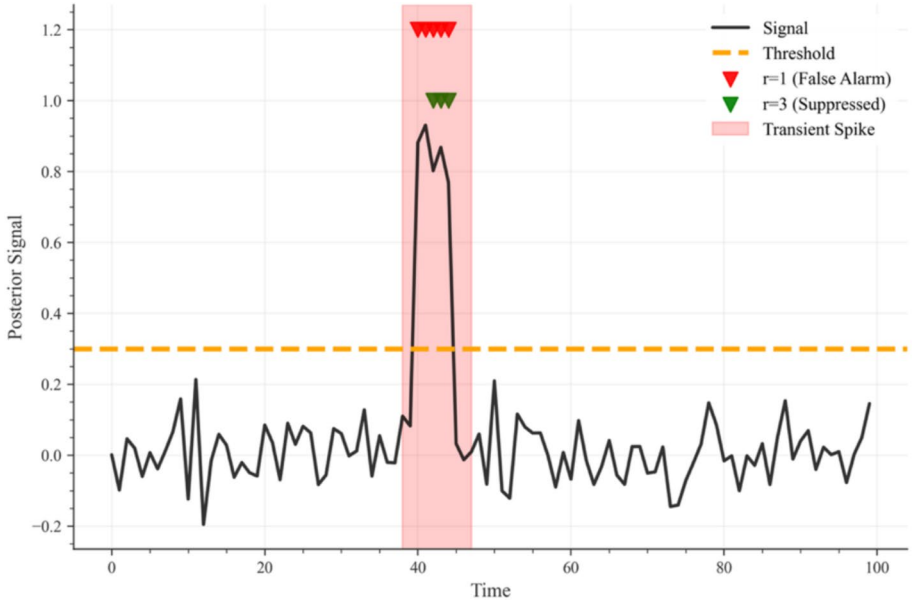

Fig. 3 Temporal trigger profiles under varying debounce runs lengths

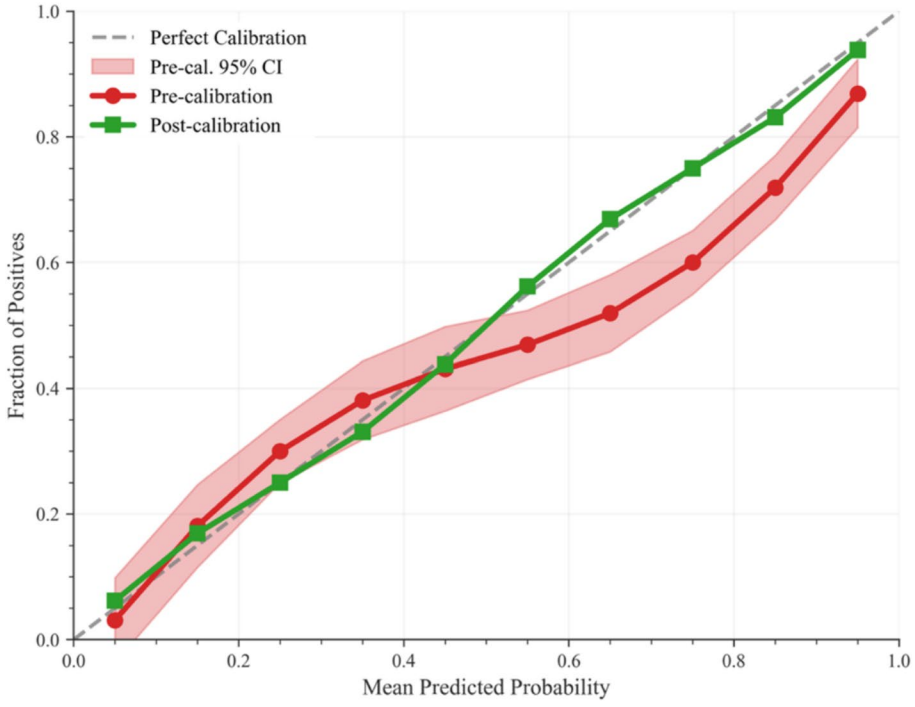

Fig. 4 Reliability diagrams of posterior drift probabilities

Table 3 Robustness under noise, missingness, and covariate shift

Scenario	Level	AUROC	AUPRC
Noise	0	0.900	0.280
	1%	0.880	0.260
Missingness	0	0.904	0.280
	10%	0.865	0.251
Shift	0	0.720	0.080
	0.5	0.930	0.340
	1.0	0.970	0.460

which is mitigated by temporal aggregation. Under covariate shift, larger departures from nominal increase detectability (AUROC $0.720 \to 0.930 \to 0.970$; AUPRC $0.080 \to 0.340 \to 0.460$), indicating that the detector benefits when the post-change distribution becomes more separable. Retuning (τ, r) and applying per-stream calibration further stabilize outcomes in noisy or sparse regimes (Table 3).

Entropy-path diagnostics corroborate the separation narrative with $\mathrm{JSD} = 0.012$, $\mathrm{KLD} = 0.657$, $\|\cdot\|_1 = 0.063$, and quadrature error 0.031, confirming the numerical stability of the boundary-aware KDE/entropy pipeline. Evidence accumulation patterns show that increasing run-length linearly suppresses isolated false positives, while late detections on weak shifts are mitigated by modestly increasing the shot budget and window size W (Fig. 5).

Paired testing against a cosine-overlap baseline using AUPRC yields a t test p value of 1.9×10^{-5} with a negative effect sign under the file's coding (effect magnitude -133.146), establishing that the observed differences are statistically reliable under the fixed-seed protocol defined in the empirical design. Holm–Bonferroni correction is applied within the comparison family.

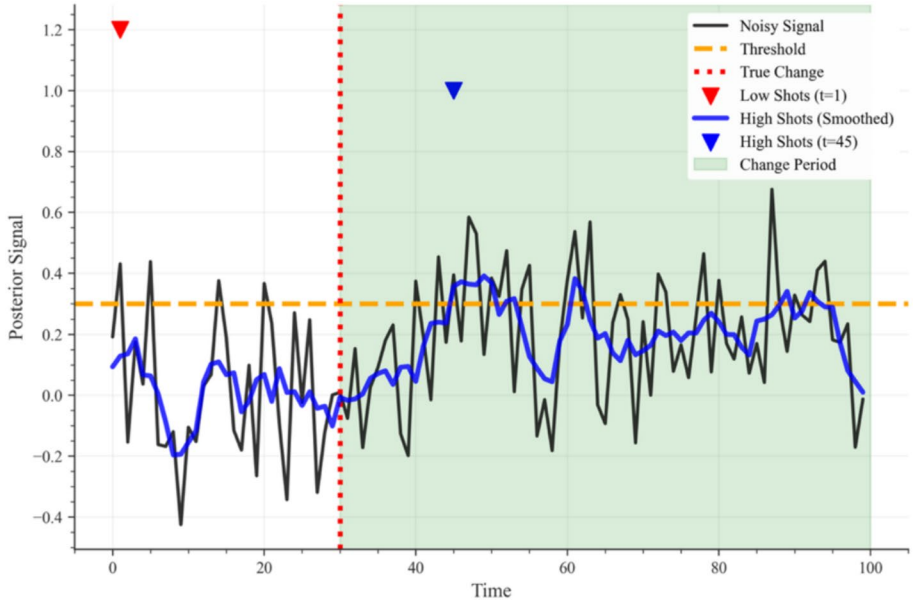

Fig. 5 Robustness diagnostics under stress scenarios

We evaluate QEK-DTMC on a real streaming anomaly dataset from NAB, as described in Sect. 4.1. The entropy trace exhibits sharp excursions at the onset of anomalous behavior (global mean of 2.823 and standard deviation 0.029), which induce a smooth posterior transition and debouncing reduces isolated threshold crossings while preserving prompt alarms under sustained deviation. Quantitatively, on this stream QEK-DTMC attains AUROC = 0.477 and AUPRC = 0.202. At the calibrated operating point used for the trace, the empirical false-alarm rate is FAR = 0.954 (alarm rate 56.45% of time steps). These results support end-to-end applicability on real data and highlight that real-stream performance is sensitive to operating-point calibration (e.g., τ, r) and preprocessing choices.

6 Discussion

We introduced QEK-DTMC, a quantum-statistical detector that couples SWAP-test overlaps with a boundary-aware kernel density, computes a Shannon entropy signal, and integrates it within a two-state DTMC posterior recursion. The design objective was to align representation, estimation, and decision so that finite-shot effects and kernel concentration are explicitly managed within a sequential rule. Empirically and theoretically, the components cohere as intended. In controlled experiments, the detector yields strong ranking performance, calibrated decision behavior, and predictable robustness under realistic stressors. Here are the main conclusions:

- Across the grid, QEK-DTMC attains AUROC 0.902 (dispersion 0.036) and AUPRC 0.284 (dispersion 0.072), with a false-alarm rate of 0.004 (dispersion 0.003) and mean detection delay of 14.0 (dispersion 6.2). Observed ranges remain tight for AUROC and AUPRC, indicating stable ranking and operating-point precision under matched budgets.
- Increasing the run-length requirement r suppresses spurious transients, approximately halving FAR when r is raised from 1 to 3, at the cost of additional

delay. Time-aligned traces show that debouncing filters isolated excursions without altering steady decisions under persistent drift.

- Posterior probabilities are reasonably calibrated out of the box, with expected calibration error 0.158 and Brier score 0.211. Mild post-hoc remapping at the operating threshold aligns nominal and empirical trigger rates and stabilizes FAR across covariate conditions.
- The detector degrades gracefully with shot-level noise and missingness, and benefits from increased separability under covariate shift. Representative trends (presented in Sect. 5.4) confirm numerical stability, with bounded quadrature error and controlled bandwidth bias.

The results suggest that choosing a ZZ feature map with sufficient expressivity to preserve task variance while avoiding collapse under finite shots yields stable kernel estimates, improved sensitivity to weak drifts, and consistent calibration across operating points. It is also desirable to maintain a moderate sliding window (e.g., ADWIN or FIFO) with conservative turnover and to set the KDE bandwidth to balance smoothing and boundary sensitivity.

Our evaluation uses synthetic streams with controlled drift mechanisms and matched classical baselines. Although this design isolates treatment effects and supports paired testing, external validity depends on application structure, hardware noise, and load regimes. The similarity stage scales linearly with window size W and realized shot count, which dominates cost relative to entropy and posterior updates. Also, although the SWAP test is simple and hardware-friendly, alternative overlap estimators can be statistically more efficient depending on device constraints. Future studies may consider semi-parametric or mixture emission models, multi-state hidden Markov model backbones, adaptive online thresholds, encoders of information-theoretic summaries with explicit expressivity-concentration trade-offs, and overlap routines that reduce estimator variance at fixed shot budgets.

Appendix

We present a complete example of the proposed QEK-DTMC detector over simple quantum states. For clarity, we show overlaps from one-qubit and two-qubit product states, and then run the end-to-end pipeline (finite-shot SWAP tests $\rightarrow$ boundary-aware KDE entropy $\rightarrow$ DTMC posterior and decision). Visual summaries are referenced where relevant: the Bloch-sphere geometry for the states, the SWAP-test circuit and shot counts, the entropy/posterior timeline, and the two-state Markov chain.

For one-qubit states, let two pure states be parameterized on the Bloch sphere by a polar angle θ (with azimuth set to zero for readability):

$$\left|\psi\left(\theta\right) rangle = \cos\frac{\theta}{2}|0rangle\right| + \sin\frac{\theta}{2}|1rangle$$

$$\left|\phi\left(\theta'\right) rangle = \cos\frac{\theta'}{2}\right| 0rangle + \sin\frac{\theta'}{2}|1rangle$$

For such pairs, the fidelity equals the squared cosine of half the geodesic angle Δ defined as $|\theta - \theta'|$ (Fig. 6):

$$s = |\langle\psi|\phi\rangle|^2 = \cos^2\frac{\Delta}{2}$$

The SWAP ancilla-0 probability is $p_0 = 1 + \frac{s}{2}$, and the unbiased similarity estimator is $\hat{s} = 2\hat{p}_0 - 1$, where $\hat{p}_0$ is the frequency of ancilla outcomes "0" across N circuit shots. Using $N = 32$, we obtain the following four one-qubit overlaps at time t against a sliding window of $W = 4$ references:

$$\begin{aligned}
\mathrm{r}_1 &: \hat{p}_0 = \tfrac{25}{32} = 0.78125 \Rightarrow \hat{s}_1 = 0.5625\,(\Delta \approx 82.82^\circ)\\
\mathrm{r}_2 &: \hat{p}_0 = \tfrac{28}{32} = 0.87500 \Rightarrow \hat{s}_2 = 0.7500\,(\Delta = 60.00^\circ)\\
\mathrm{r}_3 &: \hat{p}_0 = \tfrac{20}{32} = 0.62500 \Rightarrow \hat{s}_3 = 0.2500\,(\Delta = 120.00^\circ)\\
\mathrm{r}_4 &: \hat{p}_0 = \tfrac{30}{32} = 0.93750 \Rightarrow \hat{s}_4 = 0.8750\,(\Delta \approx 41.41^\circ)
\end{aligned}$$

Thus $S_t = \{0.5625, 0.7500, 0.2500, 0.8750\}$ (Fig. 7).

For product states $|\Psi rangle = |\,\psi_a rangle \otimes |\psi_b rangle$ and $|\Phi rangle = |\,\phi_a rangle \otimes |\phi_b rangle$ with independent one-qubit angles (Δ_a, Δ_b), the fidelity factorizes as:

$$s_{(2q)} = |langle\Psi|\,\Phi rangle|^2 = |langle\psi_a|\,\phi_a rangle|^2 \cdot |langle\psi_b|\,\phi_b rangle|^2 = \cos^2\frac{\Delta_a}{2} \cdot \cos^2\frac{\Delta_b}{2}$$

For example, choosing $\Delta_a = 60^\circ$ and $\Delta_b = 40^\circ$ gives:

$$s_{(2q)} = 0.75 \cdot \left(\cos^2 20^\circ\right) \approx 0.75 \cdot 0.883 \approx 0.662$$

Hence $p_0 \approx 0.831$. This factorization provides a convenient way to engineer target fidelities and to sanity-check SWAP counts for higher-qubit encodings of the geometry. Given the input window $W = 4$ (the $\hat{s}_i$ above), shots $N = 32$ per pair, beta-KDE concentration $\kappa = 20$, $\alpha = 0.95$, $\beta = 0.9$, $\pi_{12} = 0.20$, and decision threshold $\tau = 0.6$ (i.e., no debounce), we smooth the similarities on $[0, 1]$ with a beta-kernel mixture, and compute Shannon differential entropy $H_{13} = -0.156$. Let the observation models (calibrated on held-out pre/post-change segments) be:

$$H \,|\, S_t = 0 \sim mathcalN\,(\mu_0 = -0.85, \sigma_0 = 0.10)$$

$$H \,|\, S_t = 1 \sim mathcalN\,(\mu_1 = -0.55, \sigma_1 = 0.10)$$

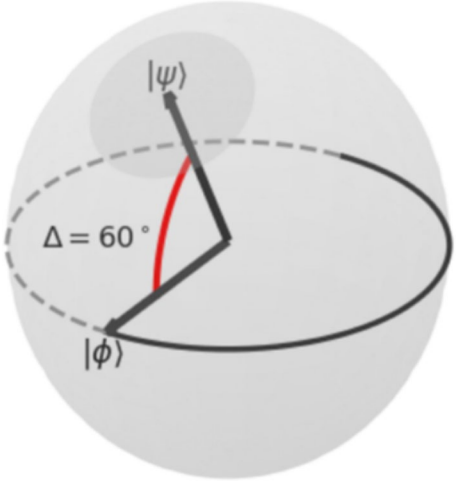

Fig. 6 Bloch-sphere illustration of encoded single-qubit states

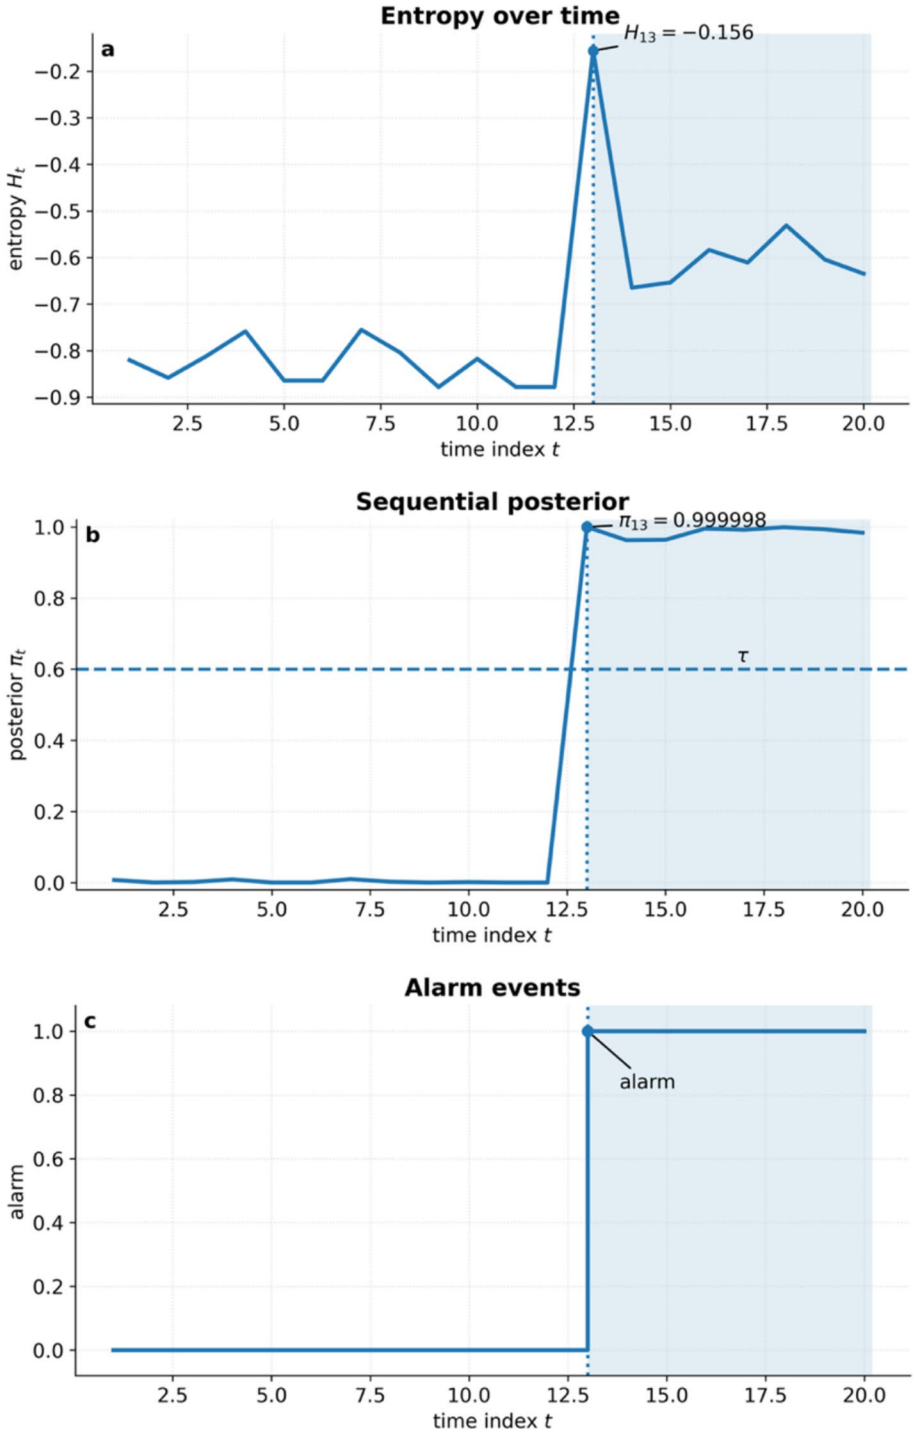

Fig. 7 Entropy-driven anomaly detection results

Compute the log-likelihood ratio at H_{13} gives $\log \frac{p(H_{13} mid S_t=0)}{p(H_{13} mid S_t=1)} = 16.324$. The predictive prior from the two-state chain is:

$$\pi_{13}^{-} = (1-\alpha)\,(1-\pi_{12}) + \beta\,\pi_{12} = 0.05 \cdot 0.8 + 0.90 \cdot 0.2 = 0.22$$

Posterior odds $O_{13}^{+} = \frac{\pi_{13}^{-}}{1-\pi_{13}^{-}}\, e^{\ell}$ and posterior $\pi_{13} = \frac{O_{13}^{+}}{1+O_{13}^{+}} = 0.9999997$. Because $\pi_{13} \geq \tau$, the detector raises an alarm at $t = 13$.

Author contributions
MK: Methodology, Writing—Review & Editing, Formal analysis, Writing- Original draft preparation, Project administration. OP: Conceptualization, Software, Data Curation, Visualization.

Funding
This research did not receive any specific grants from funding agencies in the public, commercial, or not-for-profit sectors.

Data availability
The data generated and analyzed during this study are available from the corresponding author upon reasonable request.

Declarations

Ethics and consent to participate
Not applicable.

Consent for publication
Not applicable.

Competing interests
The authors declare no competing interests.

Received: 28 September 2025 / Revised: 23 January 2026 / Accepted: 3 February 2026
Published online: 26 February 2026

References

1. Ahmad S, Lavin A, Purdy S, Agha Z. Unsupervised real-time anomaly detection for streaming data. Neurocomputing. 2017;262:134–47.
2. Armero C. Bayesian inference in Markovian queues. Queueing Syst. 1994;15:419–26. https://doi.org/10.1007/BF01189249.
3. Bravo-Prieto C. Quantum autoencoders with enhanced data encoding. Mach Learn Sci Technol. 2021;2(3):035028. https://doi.org/10.1088/2632-2153/ac0616.
4. Belis V, Woźniak KA, Puljak E, Barkoutsos P, Dissertori G, Grossi M, et al. Quantum anomaly detection in the latent space of proton collision events at the LHC. Commun Phys. 2024;7:334. https://doi.org/10.1038/s42005-024-01811-6.
5. Bennett CH, Bernstein E, Brassard G, Vazirani U. Strengths and weaknesses of quantum computing. SIAM J Comput. 1997;26(5):1510–23. https://doi.org/10.1137/S0097539796300933.
6. Bharti K, Cervera-Lierta A, Kyaw TH, Haug T, Alperin-Lea S, Anand A, et al. Noisy intermediate-scale quantum (NISQ) algorithms. Rev Mod Phys. 2022;94(1):015004. https://doi.org/10.1103/RevModPhys.94.015004.
7. Biamonte J, Wittek P, Pancotti N, Rebentrost P, Wiebe N, Lloyd S. Quantum machine learning. Nature. 2017;549(7671):195–202. https://doi.org/10.1038/nature23474.
8. Boucheron S, Lugosi G, Bousquet O. Concentration inequalities. In: Bousquet O, von Luxburg U, Rätsch G, editors. Advanced lectures on machine learning. Berlin: Springer; 2003. p. 208–40.
9. Brassard G, Høyer P, Mosca M, Tapp A. Quantum amplitude amplification and estimation. Contemp Math. 2002;305:53–74. https://doi.org/10.1090/conm/305/05215.
10. Cappé O, Moulines É, Rydén T. Inference in hidden Markov models. Berlin: Springer; 2005.
11. Cerezo M, Arrasmith A, Babbush R, Benjamin SC, Endo S, Fujii K, et al. Variational quantum algorithms. Nat Rev Phys. 2021;3:625–44. https://doi.org/10.1038/s42254-021-00348-9.
12. Cerezo M, Poremba A, Cincio Ł, Coles PJ. Variational quantum fidelity estimation. Quantum. 2020;4:248. https://doi.org/10.22331/q-2020-03-26-248.
13. Chen SX. Beta kernel estimators for density functions. Comput Stat Data Anal. 1999;31(2):131–45. https://doi.org/10.1023/A:1004165218295.
14. Cheng HC, Datta N, Liu N, Nuradha T, Salzmann R, Wilde MM. An invitation to the sample complexity of quantum hypothesis testing. NPJ Quantum Inf. 2025;11(1):94.
15. Corli S, Moro L, Dragoni D, Dispenza M, Prati E. Quantum machine learning algorithms for anomaly detection: a review. Future Gener Comput Syst. 2024;166:107632.
16. Ekert AK, Alves CM, Oi DKL, Horodecki M, Horodecki P, Kwek L-C. Direct estimations of linear and nonlinear functionals of a quantum state. Phys Rev Lett. 2002;88(21):217901. https://doi.org/10.1103/PhysRevLett.88.217901.
17. Fanizza M, Rosati M, Skotiniotis M, Calsamiglia J, Giovannetti V. Beyond the swap test: optimal estimation of quantum state overlap. Phys Rev Lett. 2020;124(6):060503. https://doi.org/10.1103/PhysRevLett.124.060503.
18. Ferrari A, Reddi A, Acharya J. Scalable nonparametric change-point detection with kernels. Pattern Recognit. 2023;138:109463. https://doi.org/10.1016/j.patcog.2023.109463.

19. García-Escartín JC, Chamorro-Posada P. SWAP test and Hong-Ou-Mandel effect are equivalent. Phys Rev A. 2013;87(5):052330. https://doi.org/10.1103/PhysRevA.87.052330.
20. Garreau D, Arlot S. Consistent change-point detection with kernels. Electron J Stat. 2018;12(2):4440–86. https://doi.org/10.1214/18-EJS1513.
21. Gasbarri G, Bilkis M, Roda-Salichs E, Calsamiglia J. Sequential hypothesis testing for continuously-monitored quantum systems. Quantum. 2024;8:1289.
22. Gretton A, Borgwardt KM, Rasch MJ, Schölkopf B, Smola AJ. A kernel two-sample test. J Mach Learn Res. 2012;13:723–73.
23. Havlíček V, Córcoles AD, Temme K, Harrow AW, Kandala A, Chow JM, et al. Supervised learning with quantum-enhanced feature spaces. Nature. 2019;567(7747):209–12. https://doi.org/10.1038/s41586-019-0980-2.
24. Hdaib M, Rajasegarar S, Pan L. Quantum deep learning-based anomaly detection for enhanced network security. Quantum Mach Intell. 2024;6(1):26.
25. Hoeffding W. Probability inequalities for sums of bounded random variables. J Am Stat Assoc. 1963;58(301):13–30.
26. Huang H-Y, Broughton M, Mohseni M, Babbush R, Boixo S, Neven H, et al. Power of data in quantum machine learning. Nat Commun. 2021;12:2631. https://doi.org/10.1038/s41467-021-22539-9.
27. Koren M, Peretz O. A quantum procedure for estimating information gain in Boolean classification task. Quantum Mach Intell. 2024;6(1):16. https://doi.org/10.1007/s42484-024-00151-6.
28. Koren M, Peretz O. Automated data-driven and stochastic imputation method. London: IntechOpen; 2024.
29. Koren M, Koren O, Peretz O. A quantum "black box" for entropy calculation. Quantum Mach Intell. 2023;5(2):37. https://doi.org/10.1007/s42484-023-00127-y.
30. Kullback S, Leibler RA. On information and sufficiency. Ann Math Stat. 1951;22(1):79–86. https://doi.org/10.1214/aoms/1177729694.
31. LaRose R, Coyle B. Robust data encodings for quantum classifiers. Phys Rev A. 2020;102(3):032420. https://doi.org/10.1103/PhysRevA.102.032420.
32. Li G, Ye R, Zhao X, Wang X. Concentration of data encoding in parameterized quantum circuits. In: Advances in neural information processing systems. Vol. 35. 2022. p. 19456–69.
33. Li J, Li Y, Song J, Zhang J, Zhang S. Quantum support vector machine for classifying noisy data. IEEE Trans Comput. 2024;73(9):2233–47.
34. Lin J. Divergence measures based on the Shannon entropy. IEEE Trans Inf Theory. 1991;37(1):145–51. https://doi.org/10.1109/18.61115.
35. Martínez Vargas E, Hirche C, Sentís G, Skotiniotis M, Carrizo M, Muñoz-Tapia R, et al. Quantum sequential hypothesis testing. Phys Rev Lett. 2021;126(18):180502.
36. McClean JR, Boixo S, Smelyanskiy VN, Babbush R, Neven H. Barren plateaus in quantum neural network training landscapes. Nat Commun. 2018;9:4812. https://doi.org/10.1038/s41467-018-07090-4.
37. Moustakides GV. Optimal stopping times for detecting changes in distributions. Ann Stat. 1986;14(4):1379–87. https://doi.org/10.1214/aos/1176350050.
38. Nielsen MA, Chuang IL. Quantum computation and quantum information (10th. anniversary. Cambridge: Cambridge University Press; 2010.
39. Page ES. Continuous inspection schemes. Biometrika. 1954;41(1–2):100–15. https://doi.org/10.1093/biomet/41.1-2.100.
40. Parker BM, Gilmour SG, Schormans JA, Maruri-Aguilar H. Optimal design of measurements on queueing systems. Queueing Syst. 2015;79(3–4):365–90. https://doi.org/10.1007/s11134-014-9421-y.
41. Peretz O, Koren M. A parameterized quantum circuit for estimating distribution measures. Quantum Machine Intelligence. 2024;6(1):22. https://doi.org/10.1007/s42484-024-00158-z.
42. Poor HV, Hadjiliadis O. Quickest detection. Cambridge: Cambridge University Press; 2008.
43. Preskill J. Quantum computing in the NISQ era and beyond. Quantum. 2018;2:79. https://doi.org/10.22331/q-2018-08-06-79.
44. Prudêncio TR, Doria SC, Rodrigues JJPC. Quantum machine learning for anomaly detection in high-energy physics. Phys Rev D. 2022;105(9):095004. https://doi.org/10.1103/PhysRevD.105.095004.
45. Robertson JK. The role of physical optics in research. Am J Phys. 1943;11:264–71. https://doi.org/10.1119/1.1990496.
46. Roldán É, Parrondo JMR. Entropy production and Kullback-Leibler divergence between stationary trajectories of discrete systems. Phys Rev E. 2012;85(3):031129. https://doi.org/10.1103/PhysRevE.85.031129.
47. Romero J, Olson JP, Aspuru-Guzik A. Quantum autoencoders for efficient compression of quantum data. Quantum Sci Technol. 2017;2:045001. https://doi.org/10.1088/2058-9565/aa8072.
48. Schuld M. Supervised quantum machine learning models are kernel methods. PRX Quantum. 2021;2(4):040101. https://doi.org/10.1103/PRXQuantum.2.040101.
49. Schuld M, Killoran N. Quantum machine learning in feature Hilbert spaces. Phys Rev Lett. 2019;122(4):040504. https://doi.org/10.1103/PhysRevLett.122.040504.
50. Shin S, Teo YS, Jeong H. Exponential data encoding for quantum supervised learning. Phys Rev A. 2023;107(1):012422. https://doi.org/10.1103/PhysRevA.107.012422.
51. Shiryaev AN. On optimum methods in quickest detection problems. Theory Probab Appl. 1963;8(1):22–46. https://doi.org/10.1137/1108002.
52. Sugiyama M, Suzuki T, Kanamori T. Density ratio estimation in machine learning. Cambridge: Cambridge University Press; 2012.
53. Swarztrauber PN. On computing the points and weights for Gauss–Legendre quadrature. SIAM J Sci Comput. 2003;24(3):945–54. https://doi.org/10.1137/S1064827500379690.
54. Tartakovsky AG, Nikiforov I, Basseville M. Sequential analysis: hypothesis testing and changepoint detection. Boca Raton, FL: Chapman & Hall/CRC; 2014.
55. Thanasilp S, Wang S, Cerezo M, Holmes Z. Exponential concentration in quantum kernel methods. Nat Commun. 2024;15:5200. https://doi.org/10.1038/s41467-024-49287-w.

56. Trefethen LN. Is Gauss quadrature better than Clenshaw-Curtis? SIAM Rev. 2008;50(1):67–87. https://doi.org/10.1137/060659831.
57. Vlasic A, Pham A. Scoring anomalous vertices through quantum walks. Ann Phys (Berl). 2025;537(5):2400282.
58. Wald A. Sequential tests of statistical hypotheses. Ann Math Stat. 1945;16(2):117–86. https://doi.org/10.1214/aoms/1177731118.
59. Ying M. Quantum computation, quantum theory and AI. Artif Intell. 2010;174(2):162–76. https://doi.org/10.1016/j.artint.2009.11.009.
60. Zong B, Jia Y, Li W, Ding Z, Zhang P, et al. A survey of quantum anomaly detection. Future Gener Comput Syst. 2024;158:107632. https://doi.org/10.1016/j.future.2024.107632.

Publisher's Note