
Kernel QuantTree

Diego Stucchi¹ Paolo Rizzo¹ Nicoló Folloni¹ Giacomo Boracchi¹

Abstract

We present Kernel QuantTree (KQT), a non-parametric change detection algorithm that monitors multivariate data through a histogram. KQT constructs a nonlinear partition of the input space that matches pre-defined target probabilities and specifically promotes compact bins adhering to the data distribution, resulting in a powerful detection algorithm. We prove two key theoretical advantages of KQT: *i*) statistics defined over the KQT histogram do not depend on the stationary data distribution ϕ_0 , so detection thresholds can be set a priori to control false positive rate, and *ii*) thanks to the kernel functions adopted, the KQT monitoring scheme is invariant to the rotation of the input data. Consequently, KQT does not require any preprocessing step like PCA. Our experiments show that KQT achieves superior detection power than non-parametric state-of-the-art change detection methods, and can reliably control the false positive rate.

1. Introduction

Change Detection (CD) is the problem of detecting distribution changes $\phi_0 \rightarrow \phi_1$ in a datastream, namely detecting when the data-generating process drifts from a stationary distribution ϕ_0 towards an unknown post-change distribution ϕ_1 . Here, we address the problem of *batch-wise CD*, where data are analyzed in fixed-size batches that, under normal conditions, contain samples drawn from ϕ_0 . The timely detection of distribution changes and the control over the false alarm rate are fundamental problems that have been widely explored in both the Machine Learning (Gama et al., 2014) and Statistical Process Control (Basville et al., 1993) literature. Among the many applications of change-detection algorithms, we mention fault detection (Tartakovsky et al., 2006), financial monitoring (Ross et al.,

2011), cryptographic attacks (Frittoli et al., 2020), and quality control (Hawkins et al., 2003).

Most CD algorithms consist of three major ingredients: *i*) a *model* $\hat{\phi}_0$ describing stationary data, which is usually learned from a training set, *ii*) a *statistical test*, where a test statistic \mathcal{T} is computed to assess the consistency of incoming data to $\hat{\phi}_0$, and *iii*) a *decision rule* on \mathcal{T} to establish whether a change has occurred. In many real-world multivariate scenarios, estimating a density model for the stationary data is often unfeasible. Therefore, non-parametric methods that describe stationary data by flexible models are preferred. Unfortunately, most non-parametric statistics are based on ranking (Ross & Adams, 2012) and can only be applied to univariate data. In Section 3, we overview the few non-parametric solutions to monitor multivariate datastreams. A relevant example is QuantTree (QT) (Boracchi et al., 2018), a change detection algorithm based on a histogram partitioning of the input space, which is supported by sound theoretical results. In particular, QT allows to operate at a controlled false alarm rate without knowing ϕ_0 nor resorting to bootstrap to estimate detection thresholds.

A fundamental limitation of QT is that splits are defined along the axis, as in Figure 1(a), resulting in a partitioning that does not always adhere to the input distribution. To mitigate this problem, a preprocessing stage is typically introduced to align the split directions to the principal components of the training set, as shown in Figure 1(b). While this procedure is often beneficial, we observe (Section 6) that it can worsen the detection performance in some unpredictable cases. Moreover, many bins in Figure 1(a)(b) have non-finite volumes, which can lead to poor estimation of bin probabilities.

In this paper, we introduce Kernel QuantTree (KQT), a non-parametric and multivariate CD algorithm that constructs histogram bins via measurable kernel functions, resulting in a powerful CD test. In contrast with the QT algorithm, which constructs bins by axis-aligned splits, KQT partitions the space in $K - 1$ compact bins defined by kernel functions evaluated on the training data. An additional bin, denoted as the *residual* bin, is non-compact and gathers all the points that do not fall in any other bin. Figure 1(c)-(d)-(e) shows that the KQT bins are compact subsets of the domain. Our intuition is that compact bins increase the flexibility when

¹Dipartimento di Elettronica, Informazione e Bioingegneria, Politecnico di Milano, Italy. Correspondence to: Diego Stucchi <diego.stucchi@polimi.it>.




Figure 1. QuantTree generates bins as intersection of hyperplanes, performing cuts along the axis (a). After a preprocessing through PCA, the cuts are oriented along the principal directions (b). Kernel QuantTree generates bins that are subsets of d -dimensional spheres according to the underlying kernel functions, namely the Euclidean (c), Mahalanobis (d) and Weighted Mahalanobis (e) distances.

modeling ϕ_0 by fitting a histogram h to training data. Moreover, estimating the bin probabilities under ϕ_0 , which are fundamental to compute the test statistic \mathcal{T}_h , is less accurate on non-compact bins.

In Section 5, we prove that KQT features two theoretical properties that have significant implications in change detection. First, the distribution of the test statistic \mathcal{T}_h computed from a KQT histogram h does not depend on the stationary distribution ϕ_0 . Consequently, detection thresholds τ can be set a priori as in QT, without knowing ϕ_0 . Second, the monitoring performed by KQT using specific kernel functions is not influenced by preprocessing based on roto-translations, including alignment to principal components. Thanks to these properties, KQT outperforms state-of-the-art alternatives on a broad experimental testbed illustrated in Section 6. In particular, KQT achieves better detection performance than the alternatives independently of preprocessing steps based on roto-translations.

In summary, these are our main contributions:

- i) We present Kernel QuantTree, a non-parametric CD method based on a histogram where bins are defined by kernel functions. KQT achieves state-of-the-art detection performance on multivariate datastreams;
- ii) We prove that statistics defined over the KQT histograms do not depend on ϕ_0 , but only on few KQT parameters. This enables control of the FPR by thresholds τ set a priori, via Monte Carlo simulations;
- iii) We prove that the monitoring performed by KQT is independent of any preprocessing by roto-translations.

2. Problem Formulation

We address the problem of change detection in batch-wise monitoring settings, where stationary data are realizations of a random vector \mathbf{X} with unknown probability density function ϕ_0 . We assume that a training set of stationary samples $\text{TR} = \{\mathbf{x}_1, \dots, \mathbf{x}_N\} \subset \mathbb{R}^d$ is provided, and that

incoming data are processed in batches W of $\nu \in \mathbb{N}$ samples each. We denote as $W \sim \phi_0$ when all the samples in the batch W are drawn from ϕ_0 .

Our goal is to design a CD algorithm that: *i*) detects distribution changes in incoming batches, and *ii*) controls the False Positive Rate (FPR), namely the probability of mistakenly detecting a change in stationary data. We formulate this CD problem as a Hypothesis Test to establish whether $W \sim \phi_0$ (null hypothesis) or $W \sim \phi_1 \neq \phi_0$, where ϕ_1 is the unknown post-change distribution. We pursue the mainstream approach of computing a test statistic \mathcal{T} on each batch W and detecting a change when

$$\mathcal{T}(W) > \tau, \quad (1)$$

where $\tau \in \mathbb{R}$ is the threshold that we set to control the FPR.

For the sake of simplicity, we assume that a batch W is either drawn from ϕ_0 or from a different unknown distribution $\phi_1 \neq \phi_0$. However, CD algorithms can in principle detect batches drawn from a mixture of ϕ_0 and ϕ_1 , even though the detection power is expected to be lower in this case.

3. Related Work

Change detection in multivariate datastreams is a challenging problem, which can be significantly simplified when ϕ_0 belongs to a known parametric family since the model $\hat{\phi}_0$ is obtained by estimating its parameters. The most popular solutions pursuing this approach consist in monitoring the likelihood of incoming data with respect to $\hat{\phi}_0$ fitted on TR. Viable options for $\hat{\phi}_0$ are Gaussian process (Saatçi et al., 2010), Gaussian Mixtures (Kuncheva, 2011) or kernel density estimators (Krempel, 2011). In (Kuncheva, 2011) and (Kuncheva & Faithfull, 2013), the Semiparametric Log-Likelihood (SPLL) algorithm fits a Gaussian Mixture Model (GMM) to TR and compares incoming batches with batches from TR by a likelihood test. Moreover, in SPLL, it is not possible to set a priori the detection threshold to control the FPR, as the distribution of the test statistic depends on ϕ_0 . Moreover, adopting a GMM to approximate ϕ_0 might not

always fit real-world data, as demonstrated by our experiments on high-dimensional datasets.

There are only a few recent multivariate methods that perform non-parametric change detection, namely, that assume that ϕ_0 and ϕ_1 are unknown. Among these, we focus on histogram-based algorithms, since these are non-parametric by design and can efficiently process datastreams in batches. As such, histogram-based algorithms represent very practical solutions to monitor multivariate datastreams. Density Tree (Criminisi et al., 2012) constructs a space partitioning by iteratively splitting regions to maximize an information-gain metric. In this case, the distribution of the test statistic depends on ϕ_0 , thus detection thresholds need to be set by bootstrap on TR. Equal Intensity K -means (EIKM) (Liu et al., 2020) divides the input space using K -means clustering, resulting in bins that yield an equal probability under ϕ_0 . EIKM is designed to handle multimodal distributions, e.g., Gaussian Mixtures, and the detection thresholds are given by asymptotic approximations of the Pearson test statistic. QuantTree (Boracchi et al., 2018) defines a partitioning \mathcal{S} of the input space in K bins by axis-aligned cuts. The theoretical properties of QT guarantee that the distribution of test statistics defined over bin probabilities does not depend on ϕ_0 , which allows to set detection thresholds a priori, with synthetically generated data through a very efficient scheme. The splits are performed such that the probability of a stationary sample to fall in each bin is close to a set of target probabilities $\{\pi_k\}$ provided as input parameters. Since the data splits are limited to the axis directions, the bins in QT require a preprocessing stage whose outcome, in terms of detection power, is uncertain, as demonstrated in our experiments (Section 6). KQT preserves the properties of QT in terms of setting detection thresholds and FPR control, and overcomes QT limitation by constructing compact bins that are not affected by roto-translations, thus better approximate the probability measure of each bin under ϕ_0 .

4. Kernel QuantTree

We present Kernel QuantTree (KQT)¹, a CD algorithm that solves a major limitation of QuantTree (QT) while generalizing and extending its theoretical guarantees. The KQT histogram is constructed by iteratively splitting the input space \mathbb{R}^d into K bins $\{S_k\}$ such that the probability of a stationary sample $\mathbf{x} \sim \phi_0$ to fall in S_k is close to a target probability π_k , which are provided as input parameters. The peculiarity of KQT is that each bin S_k for $k < K$ is defined by a *measurable kernel function* $f_k : \mathbb{R}^d \rightarrow \mathbb{R}$ and a split value $q_k \in \mathbb{R}$, and corresponds to a compact set in \mathbb{R}^d . We denote as Generalized QuantTree (GQT) partitioning the resulting histogram $h = \{(S_k, \hat{\pi}_k)\}_{k=1}^K$, which yields a partition of the input space \mathbb{R}^d , where $\hat{\pi}_k$ is the empirical

¹Code available at github.com/diegocarrera89/quantTree.

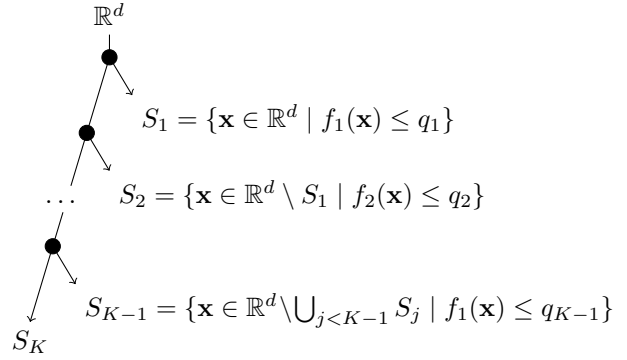


Figure 2. The Generalized QuantTree histogram is a binary splitting tree where splits isolate leaves, i.e. bins of the histogram.

probability of $\mathbf{x} \sim \phi_0$ to fall in S_k .

As illustrated in Figure 2, a GQT partitioning corresponds to an extremely imbalanced binary tree, where each split isolates a leaf, corresponding to a bin S_k . In what follows (Section 4.1 and 4.2), we illustrate in detail the GQT partitioning scheme, providing a few examples of kernel functions. In Section 5, we demonstrate that the distribution of any test statistic \mathcal{T}_h defined over a GQT partitioning does not depend on ϕ_0 , extending the theoretical results from QT. This property enables setting the detection threshold τ in (1) a priori by Monte Carlo simulations.

The monitoring scheme by KQT operates as follows. Given an input batch W containing ν test samples, we compute the test statistic $\mathcal{T}_h(W)$ and detect changes when this exceeds the threshold τ . While the theoretical properties of KQT hold for all the statistics that only depend on $\{y_k\}$, the numbers of samples in W falling in bins $\{S_k\}$, we consider the Pearson χ^2 statistic (Lehmann et al., 2005):

$$\mathcal{T}_h(W) = \mathcal{T}_h(y_1, y_2, \dots, y_K) = \sum_{k=1}^K \frac{(y_k - \nu\pi_k)^2}{\nu\pi_k}, \quad (2)$$

where $\{\pi_k\}$ are the target bin probabilities. In Section 5.3, we also prove that under some mild assumption on the kernel function f_k , this monitoring scheme becomes independent of any roto-translation applied to the data, including the PCA preprocessing. Finally, in Section 4.4, we analyze the computational complexity of KQT both at the training and monitoring stages.

4.1. Generalized QuantTree (GQT) Partitioning

The two elements defining each bin in a GQT are: *i)* a *measurable function* $f_k : \mathbb{R}^d \rightarrow \mathbb{R}$ mapping multivariate data to a single dimension and *ii)* a *split value* $q_k \in \mathbb{R}$, chosen to match the target probability π_k . Algorithm 1 illustrates the GQT histogram construction, which requires

Algorithm 1 Construction of the GQT histogram

- 1: **Input:** training set $TR = \{\mathbf{x}_i\}_{i=1}^N \subset \mathbb{R}^d$, target probabilities $\{\pi_k\}_{k=1}^K$
- 2: **Output:** GQT histogram $h = \{(S_k, \hat{\pi}_k)\}_{k=1}^K$
- 3: Set $\mathcal{X}_0 = TR$
- 4: **for** $k = 1, \dots, K - 1$ **do**
- 5: Compute $\tilde{\pi}_k = \pi_k(1 - \sum_{j < k} \pi_j)^{-1}$
- 6: Compute $\{f_k(\mathbf{x}_i)\}$ for $\mathbf{x}_i \in \mathcal{X}_{k-1}$
- 7: Set q_k as the $\tilde{\pi}_k$ -quantile of $\{f_k(\mathbf{x}_i)\}$
- 8: $S_k = \{\mathbf{x} \in \bigcap_{j < k} \overline{S}_j \mid f_k(\mathbf{x}) \leq q_k\}$
- 9: $\mathcal{X}_k = \{\mathbf{x} \in \mathcal{X}_{k-1} \mid f_k(\mathbf{x}) > q_k\}$
- 10: **end for**
- 11: $S_K = \mathbb{R}^d \setminus \bigcup_{j < K} S_j$

as input a training set TR and the target probabilities $\{\pi_k\}$. The rationale underpinning the space partitioning of KQT is to iteratively construct bins by selecting sublevel sets of the kernel functions f_k . In particular, S_1 is defined as the sublevel set of f_1 with respect to the split value q_1 :

$$S_1 = \{\mathbf{x} \in \mathbb{R}^d \mid f_1(\mathbf{x}) \leq q_1\}. \quad (3)$$

The other bins $S_k, 1 < k < K$ are obtained by isolating sublevel sets from the space that is not yet assigned to a bin

$$S_k = \{\mathbf{x} \in \bigcap_{j < k} \overline{S}_j \mid f_k(\mathbf{x}) \leq q_k\} \text{ for } k < K, \quad (4)$$

where \overline{S}_j denotes the complement of S_j in \mathbb{R}^d . The last bin, $S_K = \mathbb{R}^d \setminus \bigcup_{j < K} S_j$, is the *residual* bin, which contains all the points that do not fall in the previous bins.

The split values q_k are defined at each iteration by first identifying the set $\mathcal{X}_k \subset TR$ which contains the training points that do not fall in any bin $\{S_j \mid j < k\}$. In the beginning, no training samples have been assigned to a bin, thus, we set $\mathcal{X}_0 = TR$ and $N_0 = |TR|$ (line 3). At the k -th step, we first the percentage of points of \mathcal{X}_{k-1} that must fall in S_k to meet the target probability π_k (line 5), which we denote as:

$$\tilde{\pi}_k = \pi_k \left(1 - \sum_{j < k} \pi_j\right)^{-1}, \quad (5)$$

such that S_k contains $\pi_k N = \tilde{\pi}_k |\mathcal{X}_{k-1}|$ points. Then, we evaluate f_k on all the samples in \mathcal{X}_{k-1} (line 6), and compute the split value q_k as the $\tilde{\pi}_k$ -quantile of the projected samples $\{f_k(\mathbf{x}), \mathbf{x} \in \mathcal{X}_{k-1}\}$ (line 7). Finally, we define S_k as in (4) (line 8), and we update the set of points \mathcal{X}_k that will be used to construct the next bin (line 9). The process results in the GQT histogram $h = \{(S_k, \hat{\pi}_k)\}_{k=1}^K$, where $\hat{\pi}_k$ is the percentage of training points that fall in S_k and approximates the probability of $\mathbf{x} \sim \phi_0$ to fall in S_k .

The GQT extends the partitioning scheme underpinning QT, which corresponds to using linear split functions:

$$f_k(\mathbf{x}) = \pm 1 \cdot P_j \mathbf{x}, \quad (6)$$

where P_j is the projection over a randomly selected component j and ± 1 randomly introduces a sign flip for the projection. In the next section, we present specific measurable functions f_k that we employ in KQT.

4.2. Employed Kernel Functions in KQT

We define the kernel functions $f_k : \mathbb{R}^d \rightarrow \mathbb{R}$ as distances from a *centroid* $\mathbf{c}_k \in TR$, selected from the training set:

$$f_k(\mathbf{x}) = (\mathbf{x} - \mathbf{c}_k)^T A (\mathbf{x} - \mathbf{c}_k), \quad (7)$$

where $A \in \mathbb{R}^{d \times d}$ is the *kernel matrix*, which induces a distance measure in \mathbb{R}^d . In particular, the bins $\{S_k\}$ in (4) are subsets of d -dimensional spheres centered in $\{\mathbf{c}_k\}$ and having radii $\{q_k\}$, where the distances are measured with respect to the metric induced by A . Since spheres in \mathbb{R}^d are compact sets, all the bins S_k of a KQT, but the residual S_K , are compact and have a finite volume.

Here, we construct KQT using the Euclidean, the Mahalanobis, and the Weighted Mahalanobis (Tipping, 1999) distances, whose bins are illustrated in Figure 1. We obtain the Euclidean distance by setting $A = \mathbb{I}_d$, namely, the d -dimensional identity matrix, resulting in isotropic bins. Figure 1(c) shows that these bins poorly fit the data distribution. We obtain the Mahalanobis distance by setting $A = \Sigma^{-1}$, where $\Sigma \in \mathbb{R}^{d \times d}$ is the sample covariance matrix of TR, and in this case, the bins are anisotropic. Figure 1(d) also shows that bins are elongated towards the directions with larger variance, resulting in a better fit to the data. However, these bins poorly approximate TR when this exhibits multiple clusters, since multiple bins might span different clusters. To promote bins containing samples from a single cluster, we adopt the Weighted Mahalanobis distance and fit a Gaussian Mixture of M components to TR, and then assign a larger distance to points that belong to different components of the GMM. In KQT, we use $M = 4$ components, and the Weighted Mahalanobis kernel matrix is then defined as:

$$A(\mathbf{x}) = \frac{\sum_{m=1}^M \rho_m \cdot i_m(\mathbf{x}, \mathbf{c}) \cdot C_m^{-1}}{\sum_{m=1}^M \rho_m \cdot i_m(\mathbf{x}, \mathbf{c})}, \quad (8)$$

where μ_m , C_m , and ρ_m denote the mean, covariance matrix, and mixing probability of the m -th Gaussian, respectively. The matrix A in (8) represents a weighted average of the inverse covariance matrices of the GMM components. As in (Tipping, 1999), the weights are proportional to the mixing probabilities ρ_m and $i_m(\mathbf{x}, \mathbf{c})$, which is a computationally-tractable approximation of the distance between the point \mathbf{x} and the bin centroid \mathbf{c} .

In the following, we discuss the centroid selection strategy employed in KQT.

4.3. Centroid Selection

The criteria to select the centroids $\{\mathbf{c}_k\}$ from TR is key in KQT, as this determines both the spatial location of the bin S_k and the split value q_k associated with the kernel function f_k . Therefore, we select the centroid in TR by optimizing a partition-quality metric I , namely,

$$\mathbf{c}_k = \operatorname{argmax}_{\mathbf{c} \in \mathcal{X}_{k-1}} I[\mathbf{c}], \quad (9)$$

where \mathcal{X}_{k-1} are the training samples used to construct S_k , and $I[\mathbf{c}]$ denotes the value of the metric when we select \mathbf{c} as a centroid. In KQT, we consider two centroid selection strategies: *i*) maximizing the *information gain* associated with the split and *ii*) minimizing the *Gini index* of the distances to the centroid. For computational reasons, when \mathcal{X}_{k-1} is large, we restrict the search space to a subset of randomly sampled potential centroids $\{\mathbf{c}\} \subset \mathcal{X}_{k-1}$.

The *information gain* (Mitchell, 1997) measures the decrease in the overall entropy H after a split in the data and is typically used to assess the split quality in a data set, for example by Density Tree (Criminisi et al., 2012). The best split lowers the data entropy, maximizing the information gain. In KQT, we compute the information gain yielded by the split that divides \mathcal{X}_{k-1} into \mathcal{X}_k and $\bar{\mathcal{X}}_k = \mathcal{X}_{k-1} \setminus \mathcal{X}_k$. In particular, we can compute the entropy $H(B)$ of a set of points B using the Gaussian approximation:

$$H(B) = (1/2) \log \left((2\pi e)^d \det(\operatorname{cov}[B]) \right). \quad (10)$$

where $\operatorname{cov}[B]$ represents the sample covariance matrix computed over B . The information gain associated with the centroid \mathbf{c} is defined as

$$I[\mathbf{c}] = |\mathcal{X}_{k-1}| H(\mathcal{X}_{k-1}) - (|\bar{\mathcal{X}}_k| H(\bar{\mathcal{X}}_k) + |\mathcal{X}_k| H(\mathcal{X}_k)). \quad (11)$$

In the supplementary material, we discuss the simplifications we introduced to lower the computational burden of assessing (11) for multiple potential centroids, like the Gaussian approximation of H , which does not influence the non-parametric nature of KQT.

The *Gini index* (Gini, 1912) measures the level of uniformity in an empirical distribution and takes values between 0 (perfect equality) and 1 (maximum inequality). In KQT, we use the Gini index to prevent the selection of centroids in low-density regions. Specifically, we compute the Gini index of the distances between the training samples and the centroid as

$$I[\mathbf{c}] = \frac{\sum_{i,j} |f_k(\mathbf{x}_i) - f_k(\mathbf{x}_j)|}{2|\mathcal{X}_{k-1}| \sum_i f_k(\mathbf{x}_i)}. \quad (12)$$

We select the centroid that minimizes (12) to promote bins that cover densely populated regions of the input space.

4.4. Computational Remarks

In terms of computation cost, the training of a KQT comprises *i*) the projection of TR by f_k , whose cost depends on the specific kernel function, *ii*) the computation of the split value, which costs $O(N)$, and *iii*) the centroid selection. The cost of computing the Euclidean distance is $O(d)$, while the Mahalanobis costs $O(d^2)$ and the Weighted Mahalanobis costs $O(Md^2)$, where M is the number of Gaussian components fitted to TR. The cost of computing the information gain is dominated by the computation of the determinant in (10), which costs $O(d^3)$ while computing the Gini index only requires the distances between the training samples and the centroids, already computed to define S_k . Overall, the cost of the index computation is multiplied by the number of centroids T tested during the selection procedure by (9). Therefore, an upper bound for the cost of KQT construction is $O(KT(N + MNd^2 + d^3))$ when using the Weighted Mahalanobis distance and the information gain. During monitoring, the only operation performed is the projection by f_k of the samples of a batch W , resulting in a cost of $O(\nu KMd^2)$ in case of the Weighted Mahalanobis distance.

Table 1 reports the complexity of all the methods considered in our experiments, showing that KQT with the Weighted Mahalanobis distance is most computationally demanding, both in terms of training and inference. However, the experiments discussed in Section 6 prove that this cost is balanced by superior detection performance.

Table 1. Comparison of the computational complexity of KQT and the other considered methods, where M is the number of Gaussian components employed by KQT with the Weighted Mahalanobis distance, and R is the number of splits of Density Tree.

Method	Training Cost	Inference Cost
KQT (Weighted Maha.)	$O(KT(N + MNd^2 + d^3))$	$O(\nu KMd^2)$
QuantTree	$O(KN \log N)$	$O(\nu K)$
EIKM	$O(K^2N \log N)$	$O(\nu K)$
SPLL	$O(Nd^2)$	$O(\nu d^2)$
Density Tree	$O(KRd^3)$	$O(\nu K)$

5. Theoretical Guarantees

This section illustrates the theoretical properties of KQT and is organized as follows. In Section 5.1, we prove that the distribution of the test statistic computed by GQT over stationary data is independent of ϕ_0 , hence generalizing the main result of QT from (Boracchi et al., 2018) to a more extensive set of histogram-based monitoring schemes, including KQT. Then, in Section 5.2, we show how to exploit this result to compute detection thresholds by Monte Carlo

simulations such that the empirical FPR matches any target value α . Finally, in Section 5.3, we prove that KQT is invariant to roto-translations of the data when we use the kernel functions in Section 4.2.

5.1. Generalization of the QT Independence Theorem

The following result implies that the distribution of a test statistic like (2) computed over stationary batches by a GQT is independent of ϕ_0 , the input dimension d and the employed functions f_k . Thus, such distribution can be empirically computed via Monte Carlo simulations and used in any monitoring scenario as long as $\{\pi_k\}_k$, the training set size N , and the batch size ν are fixed.

Theorem 5.1. *Let $h = \{(S_k, \hat{\pi}_k)\}_{k=1}^K$ be a Generalized QuantTree histogram constructed using measurable functions $f_k : \mathbb{R}^d \rightarrow \mathbb{R}$, $\forall k$. Let \mathcal{T}_h be a statistic defined over batches W such that $\mathcal{T}_h(W)$ only depends on the number of samples y_1, \dots, y_K of W falling in the bins of h . Then, the distribution of \mathcal{T}_h over stationary batches $W \sim \phi_0$ depends only on the batch size ν , the number of training points N and target probabilities $\{\pi_k\}_k$.*

The proof of Theorem 5.1 follows three propositions as the proof of Theorem 1 in (Boracchi et al., 2018), which we generalize to a broader set of partitioning schemes. The first proposition states that the probability of a point drawn from the stationary distribution ϕ_0 to fall in any bin of a GQT histogram follows a Beta distribution.

Proposition 5.2. *Let $\mathbf{x}_1, \mathbf{x}_2, \dots, \mathbf{x}_M$ be i.i.d. realizations of a continuous random vector \mathbf{X} defined over $\mathcal{D} \subset \mathbb{R}^d$. Let $f : \mathbb{R}^d \rightarrow \mathbb{R}$ be a measurable function, and let $Z = f(X)$. We denote with $z_{(1)} \leq z_{(2)} \leq \dots \leq z_{(M)}$ the sorted images of $\{\mathbf{x}_j\}$ through f . For any $L \in \{1, 2, \dots, M\}$, we define the sublevel sets*

$$Q_{f,L} := \{\mathbf{x} \in \mathcal{D} : f(\mathbf{x}) \leq z_{(L)}\}. \quad (13)$$

Then, the random variable $p = P_{\mathbf{X}}(Q_{f,L})$ is distributed as Beta($L, M - L + 1$).

The proof of Proposition 5.2 is reported in the supplementary material. In the following, we denote the probability of a stationary point $\mathbf{x} \sim \phi_0$ to fall in bin k as $p_k = P_{\phi_0}(\mathbf{x} \in S_k)$. Moreover, we denote as $\tilde{p}_k = P_{\phi_0}(\mathbf{x} \in S_k \mid \mathbf{x} \notin \bigcup_{j < k} S_j)$ the probability of \mathbf{x} to fall in S_k and not in any of the previous bins.

Proposition 5.3. *For a Generalized QuantTree histogram, the following relation holds:*

$$p_k = \tilde{p}_k \cdot \left(1 - \sum_{j < k} p_j\right) = \tilde{p}_k \prod_{j < k} (1 - \tilde{p}_j). \quad (14)$$

Proposition 5.4. *For a Generalized QuantTree histogram, the random variables $\{\tilde{p}_k\}$ are independent.*

The proofs of Propositions 5.3 and 5.4 are equivalent to the proofs of the Proposition 2 and 3 for QT (Boracchi et al., 2018). Finally, the proof of Theorem 5.1 follows from Propositions 5.2-5.3-5.4 and from the following facts: *i)* the employed statistic (2) only depends on the number of samples falling in each bin $\{y_k\}$, and *ii)* when the batch is drawn from ϕ_0 , the vector $[y_1, \dots, y_K]$ is a realization of a Multinomial distribution of parameters (ν, p_1, \dots, p_K) .

5.2. Threshold Computation for FPR Control

From Theorem 5.1, it follows that in GQT we can compute a detection threshold $\tau = \tau(\alpha)$ yielding an FPR α when used as in (1) for any test statistic \mathcal{T}_h that only depends on $\{y_k\}$. For this purpose, we estimate by Monte Carlo simulations the empirical distribution of \mathcal{T}_h on stationary batches. Interestingly, (Frittoli et al., 2022) prove that the empirical distribution of N samples drawn from ϕ_0 in a QT histogram follows a Dirichlet distribution of parameters $\{\pi_1 N, \dots, \pi_{K-1} N, \pi_K N + 1\}$, where $\{\pi_k\}$ are the target probabilities used for constructing the histogram. Since the projection function does not influence the proof in (Frittoli et al., 2022), the same result holds for GQT. Moreover, the distribution of a batch $W \sim \phi_0$ in the histogram bins follows a Multinomial distribution. Thus, we can efficiently simulate the construction of a GQT and compute the values \mathcal{T}_h over $W \sim \phi_0$ by Monte Carlo simulations, such that τ is the $(1 - \alpha)$ -quantile of the resulting distribution.

Remarkably, the distribution of the test statistic does not depend on the employed kernel functions $\{f_k\}$. Therefore, we can use the same detection thresholds for any GQT that uses any measurable function f_k . Moreover, these detection thresholds also work for QT, which is a special case of GQT that uses (6).

5.3. Roto-Translation Invariance of KQT

In this section, we prove that KQT is invariant under roto-translations when the employed kernel function is either the Euclidean, Mahalanobis or Weighted Mahalanobis distance. We denote a roto-translation as $\Phi : \mathbb{R}^d \rightarrow \mathbb{R}^d$, and the image of a set $B \subset \mathbb{R}^d$ as $\Phi(B) = \{\Phi(\mathbf{x}) \mid \mathbf{x} \in B\}$. The following theorem states that the two KQT histograms $h = \{(S_k, \hat{\pi}_k)\}$ and $h' = \{(S'_k, \hat{\pi}'_k)\}$, constructed with and without preprocessing by Φ , respectively, are equivalent.

Theorem 5.5. *Let $\Phi : \mathbb{R}^d \rightarrow \mathbb{R}^d$ be a roto-translation. Let $h = \{(S_k, \hat{\pi}_k)\}$ and $h' = \{(S'_k, \hat{\pi}'_k)\}$ be the KQT histograms constructed from the training sets $\text{TR} \subset \mathbb{R}^d$ and $\text{TR}' = \Phi(\text{TR})$, where the kernel function is either the Euclidean, Mahalanobis or Weighted Mahalanobis distance. Then, we have that $S'_k = \Phi(S_k)$ and $\hat{\pi}'_k = \hat{\pi}_k$ for $k = 1 \dots, K$. In particular, for any batch W , if we compute $W' = \Phi(W)$, we have that $\mathcal{T}_h(W) = \mathcal{T}_{h'}(W')$.*

Theorem 5.5 proves that, for the considered kernel functions, the value of the test statistic computed over a batch W of data does not change when we employ a roto-translation Φ for preprocessing, including PCA, that is sometimes required for QT to achieve good detection performance. The proof of Theorem 5.5 is reported in the supplementary material and relies on the fact that the construction of the two histograms uses the same points (up to Φ), namely $\mathcal{X}'_k = \Phi(\mathcal{X}_k)$, to define each bin. In particular, we need to prove that our centroid selection strategy results in the same centroid (up to Φ) for S_k and S'_k :

Lemma 5.6. *Let \mathcal{X}_{k-1} and $\mathcal{X}'_{k-1} = \Phi(\mathcal{X}_{k-1})$ be the set of points used to construct the KQT histogram bins S_k and S'_k , respectively. Then, the centroid selection by (9) results in centroids \mathbf{c}_k and $\mathbf{c}'_k = \Phi(\mathbf{c}_k)$.*

In the supplementary material, we prove this lemma, and we use it in the proof of Theorem 5.5 to show by induction that $S_k = S'_k$ for every k . Consequently, every batch W will result in the same value of test statistic because if a point $\mathbf{x} \in W$ falls in S_k , then $\Phi(\mathbf{x}) \in W'$ will fall in S'_k .

6. Experiments

In this section, we validate KQT through several experiments, proving that KQT *i*) reaches detection performance that are statistically superior than state-of-the-art algorithms, and *ii*) can accurately control the FPR.

6.1. Datasets

We present the synthetic and real-world datasets that we employ in our experiments.

In each experiment, we consider TR made of $N = 4096$ points sampled without replacement by ϕ_0 . During testing, we randomly sample 5000 batches of ν samples from ϕ_0 and 5000 batches of ν samples from ϕ_1 to robustly assess detection performance.

Synthetic. We consider two synthetic settings with $d = 4$, the *unimodal* and the *bimodal*. In the unimodal setting, the stationary distribution ϕ_0 is a 0-mean Gaussian with a random covariance matrix. The post-change distribution ϕ_1 is obtained by roto-translation using the CCM framework (Alippi et al., 2017), such that the Kullback-Leibler distance between ϕ_0 and ϕ_1 is 1. In the bimodal setting, ϕ_0 is a Gaussian mixture of two slightly-overlapping components, and ϕ_1 is again generated by a roto-translation of each component of ϕ_0 computed using CCM. In the supplementary material, we discuss the same experiment performed with $d \in \{4, 8, 16, 32, 64, 128\}$.

INSECTS. The INSECTS dataset (Souza et al., 2020) contains feature vectors ($d = 33$) extracted from sensor measurements describing the wing-beat frequency of six (anno-

tated) species of flying insects. This dataset contains real changes caused by temperature modifications that affect the insects' flying behavior. We set up the change detection experiment such that ϕ_0 describes measurements acquired at a temperature, and the change $\phi_0 \rightarrow \phi_1$ corresponds to a temperature change. We denote as $i \rightarrow i + 1$ the considered temperature changes, with $i \in \{1, 2, 3, 4, 5\}$.

UCI. We employ real-world datasets from the UCI Machine Learning Repository (Dua & Graff, 2017) and from (Dal Pozzolo et al., 2017), with dimensions ranging from $d = 5$ to $d = 50$, reported in Table 2. We standardize these datasets and add a negligible amount of noise $\eta \sim N(0, \sigma)$ to each component to prevent the many repeated values from harming the histogram construction. The values of σ for each dataset are reported in the supplementary material. These datasets contain no distribution changes, thus stationary samples are drawn by sampling the dataset. We generate a post-change distribution ϕ_1 by shifting stationary data in a random direction with a magnitude proportional to the variance of each component.

Swarm. The Swarm Behavior classification dataset from the UCI Machine Learning Repository (Dua & Graff, 2017) comprises high-dimensional data ($d = 2400$) describing the motion of large groups of animals, which are labeled as flocking or not-flocking. We define the stationary distribution ϕ_0 as the distribution of data describing flocking groups of animals. In contrast, the post-change distribution ϕ_1 is defined by data corresponding to non-flocking animals.

High-dimensional datasets represent a challenging scenario for change detection algorithms, especially when they require estimating a density model. Therefore, these algorithms typically employ dimensionality-reduction techniques to map data to lower dimensions (Thudumu et al., 2020). To show that high-dimensional problems can be tackled by KQT upon employing such techniques, in our experiments we apply a PCA-based preprocessing step, retaining the 32 components explaining the most variance in the data.

6.2. Figures of Merit

We assess the performance of CD algorithms with two standard figures of merit, FPR and AUC. We set the detection thresholds in our experiments to yield an empirical FPR of $\alpha = 5\%$. To compare the detection power, we rank the algorithms according to their AUC, and we report their average rank (Demšar, 2006) over all the datasets and over 500 runs of each experiment. Moreover, we report the p -values of the Nemenyi post-hoc test (Nemenyi, 1963), comparing the AUCs of each method against the best-performing one. In Table 2, we mark in bold the largest AUC achieved over each dataset. We also underline values when the Nemenyi test confirms that the difference with the second best-performing

Table 2. FPR/AUC achieved by the considered methods with $K = 16$ bins and batches of $\nu = 128$ points. We report the average ranking with respect to the AUC and the p -value of the Nemenyi test. For each dataset, we mark in bold the AUCs of the best-performing method, and underline them when found to be significantly different from the second best-performing.

	d	QT (w/o PCA)	QT (w/ PCA)	KQT (Euclidean)	KQT (Mahalanobis)	KQT (Weighted Maha.)	EIKM	SPLL ($C=3$)	PCA-SPLL ($C=3$)	DT (w/o PCA)	DT (w/ PCA)
unimodal	4	4.83%/ 0.96	4.81%/ 0.98	4.86%/ 0.95	4.82%/ 0.99	4.83%/ 0.99	4.82%/ 0.87	5.46%/ 1.00	5.92%/ 0.99	7.84%/ 0.79	7.75%/ 0.81
bimodal	4	4.80%/ 0.90	4.81%/ 0.93	4.80%/ 0.90	4.81%/ 0.95	4.80%/ 0.97	4.82%/ 0.82	5.53%/ 0.92	6.02%/ 0.90	7.65%/ 0.75	7.62%/ 0.77
nino	5	5.04%/ 0.84	4.99%/ 0.91	5.00%/ 0.61	5.02%/ 0.90	5.01%/ 0.92	4.83%/ 0.53	6.14%/ 0.82	7.69%/ 0.84	7.55%/ 0.73	7.57%/ 0.58
protein	9	4.97%/ 0.90	4.98%/ 0.98	4.97%/ 0.62	4.98%/ 0.99	5.03%/ 0.99	4.88%/ 0.51	13.15%/ 0.92	8.42%/ 0.95	7.65%/ 0.70	7.64%/ 0.59
spruce	10	4.81%/ 1.00	4.83%/ 1.00	4.82%/ 0.60	4.84%/ 1.00	4.90%/ 1.00	4.86%/ 0.51	11.43%/ 1.00	11.56%/ 1.00	7.56%/ 1.00	7.57%/ 1.00
lodgepole	10	4.83%/ 1.00	4.82%/ 1.00	4.85%/ 0.65	4.80%/ 1.00	4.90%/ 1.00	4.92%/ 0.51	10.78%/ 1.00	10.89%/ 1.00	7.60%/ 1.00	7.58%/ 1.00
credit	28	4.83%/ 0.70	4.96%/ 0.87	4.89%/ 0.60	4.85%/ 0.78	5.06%/ 1.00	4.96%/ 0.51	8.67%/ 0.60	16.06%/ 0.66	7.63%/ 0.69	7.59%/ 0.82
insects (1 \rightarrow 2)	33	4.92%/ 1.00	4.93%/ 0.96	4.91%/ 0.96	4.93%/ 0.97	5.19%/ 0.99	4.93%/ 0.84	5.90%/ 0.81	6.48%/ 0.87	7.57%/ 1.00	7.60%/ 1.00
insects (2 \rightarrow 3)	33	4.93%/ 0.99	4.91%/ 1.00	4.92%/ 1.00	4.96%/ 1.00	5.25%/ 1.00	4.96%/ 0.96	5.54%/ 1.00	6.16%/ 1.00	7.60%/ 1.00	7.59%/ 1.00
insects (3 \rightarrow 4)	33	4.92%/ 0.98	4.89%/ 0.90	4.90%/ 0.90	4.88%/ 0.94	5.22%/ 0.99	4.89%/ 0.83	6.09%/ 0.75	6.69%/ 0.74	7.59%/ 1.00	7.54%/ 1.00
insects (4 \rightarrow 5)	33	4.92%/ 1.00	4.95%/ 1.00	4.91%/ 1.00	4.92%/ 1.00	5.25%/ 1.00	4.91%/ 0.95	5.48%/ 1.00	6.01%/ 1.00	7.63%/ 1.00	7.56%/ 1.00
insects (5 \rightarrow 6)	33	4.91%/ 1.00	4.90%/ 0.97	4.90%/ 0.98	4.92%/ 0.99	5.26%/ 1.00	4.90%/ 0.96	5.86%/ 0.98	6.19%/ 0.98	7.61%/ 1.00	7.63%/ 1.00
sensorless	48	4.84%/ 0.86	5.01%/ 1.00	4.82%/ 0.54	5.01%/ 1.00	7.42%/ 1.00	4.93%/ 0.50	4.33%/ 1.00	4.83%/ 1.00	7.55%/ 0.74	7.58%/ 0.60
particle	50	4.85%/ 0.89	4.87%/ 0.93	4.81%/ 0.55	4.94%/ 0.98	5.80%/ 0.99	4.84%/ 0.51	5.93%/ 0.84	6.07%/ 0.90	7.52%/ 0.80	7.60%/ 0.54
Average Ranking		5.24	4.93	7.08	3.82	2.98	9.37	5.57	5.34	5.11	5.56
Nemenyi p -value		$< 10^{-16}$	$< 10^{-16}$	$< 10^{-16}$	$< 10^{-16}$	-	$< 10^{-16}$	$< 10^{-16}$	$< 10^{-16}$	$< 10^{-16}$	$< 10^{-16}$

method is statistically significant. Confidence intervals are reported in the supplementary material.

6.3. Methods

We configure all the histogram-based methods to partition the space in K bins with uniform target probabilities $\pi_k = \frac{1}{K}$, as advised by (Boracchi et al., 2018). The number of bins K and the batch size ν must be chosen to guarantee that batches contain enough samples for a stable measure of these target probabilities. In particular, since histograms approximate the probability of a point falling in a bin by the number of training samples per batch that falls in each bin, we expect that a larger number of points per bin (i.e., the ratio ν/K) yield better detection performance. We confirm this by considering two settings: the high-ratio setting ($K = 16$, $\nu = 128$) and the low-ratio one ($K = 32$, $\nu = 64$).

QuantTree. QuantTree (Boracchi et al., 2018) uses a histogram to monitor incoming batches while controlling the FPR. Bins are constructed with axis-oriented splits along random components, and the changes are detected by the Pearson statistic. Detection thresholds are computed via Monte Carlo simulations. We test QT with and without PCA preprocessing of the data.

Equal Intensity K-Means (EIKM). EIKM (Liu et al., 2020) constructs a histogram with K bins using a K -means clustering to guarantee an equal proportion of stationary data in each bin. EIKM uses the Pearson statistic and its asymptotic approximation to set the detection thresholds.

Semiparametric Log-Likelihood (SPLL). SPLL (Kuncheva, 2011) models the stationary distribution ϕ_0 as a Gaussian Mixture Model (GMM) and suggests fitting $C = 3$ components. During inference, the test statistic associated with a batch W is computed by an upper bound of the log-likelihood of its samples. Since SPLL comes

without a threshold computation strategy, we employ Welch’s t -test (Welch, 1947) to detect batches whose average SPLL is statistically different from the training set.

PCA-SPLL. Presented in (Kuncheva & Faithfull, 2013), PCA-SPLL extends SPLL by transforming data through the PCA and monitoring by SPLL only the components with the lowest variance. Here, we keep up to 5% of the variance and fit $C = 3$ Gaussians.

Density Tree. Inspired by (Criminisi et al., 2012), the bins of Density Tree minimize the entropy after each split. Density Tree employs the Pearson test statistic, and sets the detection threshold via bootstrapping over a portion of training data. We test Density Tree with and without PCA preprocessing of the data.

6.4. Results and Discussion

Table 2 reports the FPR and AUC achieved by all the considered methods in the high-ratio setting, averaged over 500 runs. Here, we only report the performance of KQT maximizing the information gain (11) to select centroids. In the supplementary material, we also report the results for the low-ratio and a comparison showing that minimizing the Gini index (12) leads to comparable performance.

Table 3. FPR and AUC achieved by QuantTree and Kernel QuantTree on the Swarm dataset processed by a PCA to retain $d = 32$ components. In parenthesis, the standard deviation of the results.

	FPR	AUC
QuantTree	4.61% (1.65%)	1.00 (0.00)
KQT (Euclidean)	4.53% (1.63%)	1.00 (0.00)
KQT (Mahalanobis)	4.50% (1.58%)	1.00 (0.00)
KQT (Weighted Maha.)	4.61% (1.67%)	1.00 (0.00)

In most experiments, the empirical FPR achieved by KQT is

close to the target $\alpha = 5\%$ for all the considered kernel functions. However, the KQT with the Weighted Mahalanobis distance struggles to control the FPR when d is large. This is a known limitation of high-dimensional settings, where the estimated GMM might be poorly conditioned when TR is not sufficiently large. Therefore, when the GMM fit from TR yields Gaussians having covariances with large condition numbers, it is convenient to use KQT with the Mahalanobis distance. This latter, in fact, can control the FPR and usually achieves comparable detection performance with the Weighted Mahalanobis distance.

To further investigate how the dimension d influences the detection performance of KQT, we run another experiment where we train KQT on synthetic data with $d \in \{4, 8, 16, 32, 64, 128\}$ and with $N \in \{4096, 16384\}$. The results reported in the supplementary material prove that the FPR control worsens when d increases, but also that using large training sets heavily mitigates this problem. Moreover, this issue can be avoided by employing dimensionality reduction techniques on high-dimensional data. To this purpose, we run an experiment on the Swarm dataset ($d = 2400$) preprocessed by a PCA transformation that retains only the first 32 principal components. Table 3 shows that thanks to this preprocessing, KQT can seamlessly operate without incurring the loss of FPR control observed on some UCI datasets. The large AUC achieved by all the methods on this dataset proves that classes are very far apart, and we argue that a similar situation would have happened if we had artificially added a change (which would not correspond to a real-world problem) to offset each component. In fact, even a small perturbation would result in a very apparent change. This is probably the reason why change detection benchmarks are of lower dimensions (e.g., INSECTS $d=33$).

As for the other methods, QT and EIKM accurately control the FPR. In contrast, SPLL and PCA-SPLL mostly exceed the target, and we speculate that the distribution of the SPLL test statistic does not satisfy the t -test assumptions. Finally, Density Tree largely overshoots the target FPR, being unable to learn a detection threshold from bootstrapping.

KQT with the Weighted Mahalanobis and the Mahalanobis distance represent the best and second-best method in terms of AUC, mostly outperforming the alternatives in most settings. In particular, the advantage over the third-ranked method (QT w/ PCA) is considerable, and the p -values of the Nemenyi post-hoc test show that the advantage of KQT is also statistically significant. However, when using the Euclidean distance on real-world data, the performance of KQT worsens because its anisotropic bins cannot model the intricate data distributions of real data. As for the other methods, SPLL performs well over the synthetic datasets but fails over the INSECTS and UCI, achieving significantly low performance on the latter. Instead, PCA-SPLL mainly

improves the performance of SPLL, even though it cannot compete with the top-performing methods. Finally, Density Tree mostly achieves low detection performance, except on the INSECTS dataset, where it surpasses the other methods.

Our experiments show that preprocessing by PCA is in general beneficial for QT, as the average rank of QT w/PCA is lower than QT w/o PCA. However, in some settings, QT w/o PCA performs better. In contrast, KQT achieves the best AUC independently of the PCA preprocessing, thanks to the invariance to roto-translation proved in Section 5.3, and Density Tree is also not affected by the PCA.

The supplementary material reports the detection results in the low-ratio setting ($K = 32, \nu = 64$). Overall, the results conform to those in the high-ratio setting and confirm that a larger expected number of points per bin improves the detection performance for all the methods. For the same reason, QT and KQT achieve lower FPR in the low-ratio setting than in the high-ratio, since the Pearson statistic assumes fewer distinct values. Thus, while still controlling the FPR, this results in a slightly lower percentage of false alarms. The KQT with the Weighted Mahalanobis distance achieves the highest AUC in the low-ratio setting, with a statistically significant advantage over the competitors.

We conclude by remarking that data in the INSECTS and UCI datasets are not drawn from multivariate Gaussian distributions, as suggested by the low performance achieved by SPLL, which is based on a GMM. To confirm this, we run the Shapiro-Wilk normality test on the marginals of our real-world data, showing that these are not univariate Gaussians. In the supplementary material, we report the p -values of these tests, which are in the range of 10^{-20} .

7. Conclusions

In this paper we presented KQT, a non-parametric multivariate change detection method for batch-wise monitoring. KQT constructs a space partitioning via kernel functions, resulting in bins which are compact and lead to superior detection power. We compare our method to several state-of-the-art approaches for CD, achieving the best results in terms of detection power and false positives control. Future work includes integrating KQT in a sequential monitoring scheme, where data are processed in a continuous stream.

8. Acknowledgments

This paper is supported by the PNRR-PE-AI FAIR project funded by the NextGeneration EU program.

References

- Alippi, C., Boracchi, G., and Carrera, D. Ccm: Controlling the change magnitude in high dimensional data. In Angelov, P., Manolopoulos, Y., Iliadis, L., Roy, A., and Vellasco, M. (eds.), *Advances in Big Data*, pp. 216–225, Cham, 2017. Springer International Publishing. ISBN 978-3-319-47898-2.
- Basseville, M., Nikiforov, I. V., et al. *Detection of abrupt changes: theory and application*, volume 104. prentice Hall Englewood Cliffs, 1993.
- Boracchi, G., Carrera, D., Cervellera, C., and Maccio, D. Quantree: Histograms for change detection in multivariate data streams. In *International Conference on Machine Learning*, pp. 639–648. PMLR, 2018.
- Criminisi, A., Shotton, J., and Konukoglu, E. Decision forests: A unified framework for classification, regression, density estimation, manifold learning and semi-supervised learning. *Foundations and trends® in computer graphics and vision*, 7(2–3):81–227, 2012.
- Dal Pozzolo, A., Boracchi, G., Caelen, O., Alippi, C., and Bontempi, G. Credit card fraud detection: a realistic modeling and a novel learning strategy. *IEEE transactions on neural networks and learning systems*, 29(8):3784–3797, 2017.
- Demšar, J. Statistical comparisons of classifiers over multiple data sets. *The Journal of Machine Learning Research*, 7:1–30, 2006.
- Dua, D. and Graff, C. UCI machine learning repository, 2017. URL <http://archive.ics.uci.edu/ml>.
- Frittoli, L., Matteo, B., Silvia, M., Carrera, D., Beatrice, R., Fragneto, P., Ruggero, S., and Boracchi, G. Strengthening sequential side-channel attacks through change detection. In *Conference on Cryptographic Hardware and Embedded Systems (CHES)*, volume 3, pp. 1–21, 2020.
- Frittoli, L., Carrera, D., and Boracchi, G. Nonparametric and online change detection in multivariate datastreams using quantree. *IEEE Transactions on Knowledge and Data Engineering*, 2022.
- Gama, J., Žliobaitė, I., Bifet, A., Pechenizkiy, M., and Bouchachia, A. A survey on concept drift adaptation. *ACM computing surveys (CSUR)*, 46(4):1–37, 2014.
- Gini, C. *Variabilità e mutabilità: contributo allo studio delle distribuzioni e delle relazioni statistiche.[Fasc. I.]*. Tipogr. di P. Cuppini, 1912.
- Hawkins, D. M., Qiu, P., and Kang, C. W. The changepoint model for statistical process control. *Journal of quality technology*, 35(4):355, 2003.
- Kreml, G. The algorithm apt to classify in concurrence of latency and drift. In *Proceedings of the Intelligent Data Analysis (IDA)*, pp. 222–233, 2011.
- Kuncheva, L. I. Change detection in streaming multivariate data using likelihood detectors. *IEEE transactions on knowledge and data engineering*, 25(5):1175–1180, 2011.
- Kuncheva, L. I. and Faithfull, W. J. Pca feature extraction for change detection in multidimensional unlabeled data. *IEEE transactions on neural networks and learning systems*, 25(1):69–80, 2013.
- Lehmann, E. L., Romano, J. P., and Casella, G. *Testing statistical hypotheses*, volume 3. Springer, 2005.
- Liu, A., Lu, J., and Zhang, G. Concept drift detection via equal intensity k-means space partitioning. *IEEE transactions on cybernetics*, 51(6):3198–3211, 2020.
- Mitchell, T. M. *Machine learning*, volume 1. McGraw-hill New York, 1997.
- Nemenyi, P. B. *Distribution-free multiple comparisons*. Princeton University, 1963.
- Ross, G. J. and Adams, N. M. Two nonparametric control charts for detecting arbitrary distribution changes. *Journal of Quality Technology*, 44(2):102, 2012.
- Ross, G. J., Tasoulis, D. K., and Adams, N. M. Nonparametric monitoring of data streams for changes in location and scale. *Technometrics*, 53(4):379–389, 2011.
- Saatçi, Y., Turner, R. D., and Rasmussen, C. E. Gaussian process change point models. In *Proceedings of the International Conference on Machine Learning (ICML)*, pp. 927–934, 2010.
- Souza, V. M., dos Reis, D. M., Maletzke, A. G., and Batista, G. E. Challenges in benchmarking stream learning algorithms with real-world data. *Data Mining and Knowledge Discovery*, 34(6):1805–1858, 2020.
- Tartakovsky, A. G., Rozovskii, B. L., Blazek, R. B., and Kim, H. A novel approach to detection of intrusions in computer networks via adaptive sequential and batch-sequential change-point detection methods. *IEEE TSP*, 54(9):3372–3382, 2006.
- Thudumu, S., Branch, P., Jin, J., and Singh, J. A comprehensive survey of anomaly detection techniques for high dimensional big data. *Journal of Big Data*, 7:1–30, 2020.
- Tipping, M. E. Deriving cluster analytic distance functions from gaussian mixture models. 1999.
- Welch, B. L. The generalization of ‘student’s’ problem when several different population variances are involved. *Biometrika*, 34(1-2):28–35, 1947.