



Estimating Tukey depth using incremental quantile estimators

Hugo L. Hammer^{a,*}, Anis Yazidi^b, Håvard Rue^c

^aOsloMet - Oslo Metropolitan University, Norway Simula Metropolitan Centre for Digital Engineering, Norway

^bOsloMet - Oslo Metropolitan University, Norway

^cKing Abdullah University of Science and Technology, Saudi Arabia



ARTICLE INFO

Article history:

Received 26 August 2020

Revised 13 September 2021

Accepted 18 September 2021

Available online 20 September 2021

Keywords:

Data stream

Incremental quantile estimator

Distributional patterns

Real-time analytics

Tukey depth

ABSTRACT

Measures of distance or how data points are positioned relative to each other are fundamental in pattern recognition. The concept of depth measures how deep an arbitrary point is positioned in a dataset, and is an interesting concept in this regard. However, while this concept has received a lot of attention in the statistical literature, its application within pattern recognition is still limited.

To increase the applicability of the depth concept in pattern recognition, we address the well-known computational challenges associated with the depth concept, by suggesting to estimate depth using *incremental quantile estimators*. The suggested algorithm can not only estimate depth when the dataset is known in advance, but can also *track* depth for dynamically varying data streams by using *recursive updates*. The tracking ability of the algorithm was demonstrated based on a real-life application associated with detecting changes in human activity from real-time accelerometer observations. Given the flexibility of the suggested approach, it can detect virtually *any* kind of changes in the distributional patterns of the observations, and thus outperforms detection approaches based on the Mahalanobis distance.

© 2021 The Authors. Published by Elsevier Ltd.

This is an open access article under the CC BY license (<http://creativecommons.org/licenses/by/4.0/>)

1. Introduction

Measures of distance or how data points are positioned relative to each other, are fundamental in pattern recognition. For example in anomaly detection [9] e.g using auto encoders [4,35], clustering [13,25] or classification [17,29].

To measure distance or how data points are positioned relative to each other, *data depth* is an interesting concept. Data depth measures how deep an arbitrary point is position in a dataset. While the concept has received a lot of attention in the statistical literature [27], the application within pattern recognition is still limited. There are however some notable exceptions. For example [14,18,19] applied the concept for classification and clustering. Depth has also been applied to a wide range of disciplines such as economy [14,19,21], health and biology [15,33], ecology [3] and hydrology [5] to name a few.

The earliest and most popular depth measure is Tukey depth [32]. The Tukey depth of a point is defined as the minimum probability mass carried by any closed halfspace containing the point. However, the computation of Tukey depth for higher dimensions or

for even moderate amounts of data is computationally demanding which limits its applicability [23].

The main aim of this paper is to address these aforementioned computational issues, and thus increase the applicability of the depth concept within pattern recognition. Our approach takes advantage of the following result from Kong and Mizera [20] according to which the authors defined halfspaces such that a specific portion of the data points are on one side of the halfspace. They further showed that contours with a specific Tukey depth can be estimated from the intersection of such halfspaces over different directions. Such contours can again be used to estimate the depth of any point. In order to apply this result to estimate depth in dimension p , the positions of $O(c^{p-1})$, $c > 1$ halfspaces must be estimated requiring estimators that are both memory and computationally efficient. In this paper, we therefore suggest to use *incremental quantile estimators* to estimate the positions of the halfspaces [10,11]. These estimators only need to store a single value in memory, i.e. $O(1)$, and only need to perform a single operation per observation resulting in a computational complexity of $O(n)$ for n observations. As opposed to this, traditional quantile estimators have a memory requirement of $O(n)$ and a $O(n \log n)$ computational complexity. However, the computational efficiency comes with a price and traditional estimators provide more precise estimates based on the same observations.

* Corresponding author.

E-mail address: hugo.hammer@oslomet.no (H.L. Hammer).




Fig. 1. Examples of three halfspaces (blue, green, red) containing the point x . The blue contours represent some probability distribution P . (For interpretation of the references to colour in this figure legend, the reader is referred to the web version of this article.)

The second aim is to recursively update and even track Tukey depth contours of streams of multivariate data in real time. A remarkable advantage with incremental quantile estimators is that they not only can estimate quantiles when the data is known in advance, but can recursively update and even track quantiles of data streams. Thus, by using incremental quantile estimators to estimate halfspaces, Tukey depth contours can be tracked in real time. We are not aware of any other method that can efficiently compute Tukey depth in real-time.

Finally, the real-world applicability of computing Tukey depth in real-time settings is demonstrated where the developed methods are used to detect changes in human activity in real-time from accelerometer observations. Due the flexibility of the suggested approach, it can detect virtually *any* kind of changes in the distributional patterns of the accelerometer observation, and outperforms popular approaches based on Mahalanobis distance.

The main contributions of the paper are as follows:

- We present a new, simple and computationally efficient method to compute Tukey depth.
- The method can even be used to compute Tukey depth in real-time, and is to the best of our knowledge the first method with this ability.
- The method is applied to detect changes in human activity in real-time which demonstrates its usefulness and applicability in real-world scenarios.

The paper is organized as follows. In [Section 2](#), the concept of depth is introduced including some theoretical fundamentals to compute Tukey depth. [Section 3](#) provides an efficient procedure to estimate Tukey depth. [Section 4](#) presents performance metrics that will be used to evaluate the algorithm and [Sections 5](#) and [6](#) provide synthetic and real-life data experiments.

2. The concept of depth

Let $X = (X_1, \dots, X_p)^T$ represent a p -dimensional stochastic vector with probability distribution P . Let $D(x, P)$ denote the *depth function* of a point x with respect to the probability distribution P . A high (low) value of the depth function refers to a central (outlying) point of the probability distribution. A general depth function is defined by satisfying the natural requirements of affine invariance, maximality at center, monotonicity relative to deepest point and vanishing at infinity [36].

The most used depth function is Tukey depth.

Definition 1 (Tukey depth). Let \mathcal{U} refer to the set of all vectors with unit length. Tukey depth is the minimum probability mass carried by any closed halfspace containing the point

$$D(x, P) = \inf_{u \in \mathcal{U}} P(u^T X \leq u^T x) \quad (1)$$

[Fig. 1](#) shows three halfspaces containing the point x . The probability mass carried by the red and blue halfspaces are larger than

for the green halfspace. Thus the probability mass carried by the green halfspace is closer to the Tukey depth, which is the minimum probability mass over all half spaces containing x . Intuitively, this is a reasonable and general measure for the centrality of x with respect to P .

Define α -depth region, directional quantile and directional quantile halfspace.

Definition 2 (α -depth region). The α -depth region with respect to Tukey depth, $D(\alpha)$, is defined as the set of points whose depth is at least α

$$D(\alpha) = \{x \in \mathbb{R}^p : D(x, P) \geq \alpha\} \quad (2)$$

The boundary of $D(\alpha)$ is known as the α -depth contour.

The α -depth regions are closed, convex, and nested for increasing α .

Definition 3 (Directional quantile). For any unit directional vector $u \in \mathcal{U}$, define the *directional quantile* as

$$Q(\alpha, u^T X) = F_{u^T X}^{-1}(\alpha) \quad (3)$$

where $F_{u^T X}^{-1}(x)$ refers to the inverse of the univariate cumulative distribution function of the projection of X on u .

Definition 4 (Directional quantile halfspace). The directional quantile halfspace is defined as

$$H(\alpha, u) = \{x \in \mathbb{R}^p : u^T x \geq Q(\alpha, u^T X)\} \quad (4)$$

which is bounded away from the origin at distance $Q(\alpha, u^T X)$ by the hyperplane with normal vector u .

Consequently $P(X \in H(\alpha, u)) = 1 - \alpha$ for any $u \in \mathcal{U}$.

The estimation procedures in this paper builds on the following theorem from [20].

Theorem 1. The α -depth region in (2) equals the directional quantile envelope

$$D(\alpha) = \bigcap_{u \in \mathcal{U}} H(\alpha, u) \quad (5)$$

Tukey depth may not be defined for depths above some threshold and the intersection becomes empty.

3. Efficient estimation of tukey depth

Given a multivariate dataset, in this section we suggest a simple procedure to estimate whether an arbitrary point is within or outside an α -depth region. The procedure uses [Theorem 1](#), and consists of three parts.

1. Unit length directional vectors. The generation of uniformly distributed directional vectors is simple: Let Z_1, \dots, Z_p be independent standard normally distributed stochastic variables and define $Z = (Z_1, \dots, Z_p)^T$. Then $U = Z/\|Z\|_2$ will be uniformly distributed on the unit sphere, where $\|\cdot\|_2$ refers to the Euclidean

norm. This is the procedure used in most of the experiments in this paper. However, intuitively, using directional vectors that are more equidistantly spread on the unit sphere would be more efficient. We thus also considered the following approach according to which we generate many uniformly distributed directional vectors, N_u , and secondly filter out directional vectors that are closer than some threshold. The approach is however computationally demanding, $O(N_u^2 p^2)$. There are other algorithms to generate fairly equidistantly spread direction vectors, see e.g spiral algorithms [30]. We have not evaluated the potential of these algorithms.

2. Directional quantiles estimates. The next part is to estimate directional quantiles for each directional vector generated above. As pointed out in the introduction, we use incremental quantile estimators. A prominent example is the DUMIQE algorithm which recursively updates directional quantile estimates as follows for every observation $u_i^T x_{j-1}$ [34]

$$\begin{aligned} \widehat{Q}(\alpha, u_i^T X_j) &\leftarrow (1 + \lambda\alpha)\widehat{Q}(\alpha, u_i^T X_{j-1}), & \text{if } u_i^T x_j > \widehat{Q}(\alpha, u_i^T X_{j-1}) \\ \widehat{Q}(\alpha, u_i^T X_j) &\leftarrow (1 - \lambda(1 - \alpha))\widehat{Q}(\alpha, u_i^T X_{j-1}), & \text{if } u_i^T x_j < \widehat{Q}(\alpha, u_i^T X_{j-1}) \end{aligned} \quad (6)$$

The update is quite intuitive. If the sample $u_i^T x_j$ is above (respectively below) the current estimate, increase (respectively reduce) the corresponding directional quantile estimate. The tuning parameter $\lambda > 0$ controls the update size. If the data is known beforehand or it comes in the form of a stationary data stream, it makes sense to let the value of λ be reduced with time. For non-stationary data streams a constant value of λ is more suitable to gradually forget old and outdated data [34]. The procedure is detailed in Algorithm 1.

Algorithm 1 Estimating directional quantiles.

Input:

u_1, \dots, u_{n_u} // Unit length directional vectors
 x_1, x_2, \dots, x_n // Dataset
 α, λ
 $\widehat{Q}(\alpha, u_i^T X_0)$ // Initial value

Method:

```

1: for  $j \in 1, 2, \dots, n$  do
2:   for  $i \in 1, 2, \dots, n_u$  do
3:     if  $u_i^T x_j > \widehat{Q}(\alpha, u_i^T X_{j-1})$  then
4:        $\widehat{Q}(\alpha, u_i^T X_j) \leftarrow (1 + \lambda\alpha)\widehat{Q}(\alpha, u_i^T X_{j-1})$ 
5:     else
6:        $\widehat{Q}(\alpha, u_i^T X_j) \leftarrow (1 - \lambda(1 - \alpha))\widehat{Q}(\alpha, u_i^T X_{j-1})$ 
7:     end if
8:   end for
9: end for

```

3. Compute if a point w is within the α -depth region. The directional quantile estimates from Algorithm 1 can be used to compute if w is within all the directional quantile halfspaces and thus, according to Theorem 1, being within the α -depth region. The procedure is detailed in Algorithm 2. The condition in line 2 is based on Equation (4) and checks if w is outside of the estimated directional quantile halfspace.

Convergence. The procedure consists of two approximations 1) the finite number of directional vectors and 2) the estimates of the true directional quantiles. To ensure convergence, the directional vector selection procedure must cover the unit sphere when the number of directional vectors goes to infinity and, secondly, the directional quantile estimates must converge to the true directional quantiles, when the number of observations goes to infinity. By using the simple procedure above to select uniformly distributed directional vectors, the first requirement is satisfied. Further, [34] and [11] prove the second requirement.

Algorithm 2 Compute if a point w is within the α -depth region.

Input:

$\widehat{Q}(\alpha, u_i^T X_n), i = 1, \dots, n_u$ // Dir. quantile estimates from Algorithm 1.

$w, i = 1$

InAlphaDepthRegion = True // True (False) if w is within (outside) the α -depth region

Method:

```

1: while InAlphaDepthRegion and  $i \leq n_u$  do
2:   if  $u_i^T w < \widehat{Q}(\alpha, u_i^T X_n)$  then
3:     InAlphaDepthRegion = False
4:   end if
5:    $i \leftarrow i + 1$ 
6: end while
7: Print("Point  $w$  in  $\alpha$ -depth region?", InAlphaDepthRegion)

```




Fig. 2. Figure illustrating the approach to measure α -depth contour estimation error. The black and blue curves show the true α -depth regions and the envelope estimate. The lines with directions $v_i, i = 1, \dots, n_v$ are shown in red. (For interpretation of the references to colour in this figure legend, the reader is referred to the web version of this article.)

4. Performance metrics

We suggest to measure error along lines $l_i, i = 1, \dots, n_v$ going through the center of the true distribution and outward in uniformly distributed directions $v_i, i = 1, \dots, n_v$ (Fig. 2). This approach scales well with dimension p .

We suggest two error measures:

Depth error: Let $\tilde{w}_{i,k}$ denote the point of intercept between the line, l_i , and the envelope and compute the true depth at this point, $D(\tilde{w}_{i,k}, P)$. The error is computed using mean absolute depth error (MADE) over all the lines $l_i, i = 1, \dots, n_v$

$$\text{MADE}_k = \frac{1}{n_v} \sum_{i=1}^{n_v} |\alpha_k - D(\tilde{w}_{i,k}, P)|$$

and again average over envelopes

$$\text{MADE} = \frac{1}{K} \sum_{k=1}^K \text{MADE}_k \quad (7)$$

To compute MADE for higher dimensions, the true depth must be computed for a large set of points $\tilde{w}_{i,k}$. For non-elliptic distributions this is computationally demanding and was limited to $p \leq 6$ in the experiments. For elliptic distributions, and in particular multivariate normal distributions, the true depth of any point can be computed analytically and thus MADE was computed up to dimension $p = 10$ in the experiments. Details are given in supplementary material S.1. Obviously, if we knew that the observations were mul-

tivariate normally distributed, other depth measures such as Mahalanobis depth would be more natural, but the computations are only used to evaluate the performance of the algorithm for high dimensions.

Euclidean distance: Along each line, l_i , compute the point of intercept between the line and the true α -depth contour of depth α_k , denoted $w_{i,k}$. Compute the error as the average Euclidean distance (ED)

$$ED_k = \frac{1}{n_p} \sum_{i=1}^{n_p} \|w_{i,k} - \tilde{w}_{i,k}\|_2$$

where $\tilde{w}_{i,k}$ still refers to the intercept between line l_i and the envelope. Further take the average over envelopes

$$ED = \frac{1}{K} \sum_{k=1}^K ED_k \quad (8)$$

5. Synthetic experiments

In this section, the performance of the algorithms in Section 3 are evaluated in several synthetic experiments. The experiments focus on streaming data, except in Section 5.2. In Section 6, the algorithms are demonstrated in a real-life data example.

All computations were run on a Dell PowerEdge R815 server with 64 1.8GHz AMD CPU processors and Linux Ubuntu operating system version 16.04. The experiments were implemented in R [28], but with the most computer intensive parts in C++ integrated using Rcpp [7,8].

5.1. Synthetic experiments - Static data stream

Figs. 3 show results of estimating the $\alpha = 0.1$ depth contour for a multivariate normally distributed data stream with parameters

$$\mu = \begin{bmatrix} 0 \\ 0 \end{bmatrix}, \quad \Sigma = \begin{bmatrix} 1 & 0.82 \\ 0.82 & 1 \end{bmatrix} \quad (9)$$

Directional quantiles were estimated using DUMIQE with decreasing values of the tuning parameter, $\lambda_n = 1/n$.

We see that a fairly good estimate is achieved with 200 observations and that the error is minimal with 2000 observations. A similar visualization for the highly non-elliptical and heavy tailed lognormal distribution is shown in Figure 7 in supplementary material S.2. Due to the flexibility of the depth concept, the method performs equally well for such a distribution. Further, in supplementary material S.2, a few examples of joint estimation of multiple α -depth regions using the ShiftQ algorithm are shown. The results show that multiple depth regions can efficiently be estimated for both Gaussian and non-Gaussian distributions.

Considered now joint estimation of α -depth regions for $\alpha = 0.05, 0.2$ and 0.4 and for $p > 2$. Table 1 shows results for standard multivariate normally distributed observation. More detailed results are given in Figures 12 and 13 in supplementary material S.2. CPU time refers to the computational time needed per α -depth region to obtain estimates with a given precision using a single CPU core.

The number of directional vectors (and thus CPU time) increases with p and estimation precision. The algorithm performs very well. For example, for dimension $p = 10$, MADE less than 0.02 is obtained in about 1.5 seconds. MADE < 0.01 could be reached in shorter CPU time than what is shown in Table 1 using a higher number of directional vectors, but this is not explored.

Now, assume that $X = (X_1, \dots, X_p)^T$ is a multivariate normally distributed variable with zero expectation vector and strong dependencies

$$\text{Cov}(X_i, X_j) = \exp(-0.2|i - j|), \quad i, j = 1, \dots, p \quad (10)$$

Table 1

Multivariate standard normal distribution case: The second and third columns show the CPU time (in seconds) and the number of directional vectors used to obtain MADE less than 0.05. The other columns show the same to obtain MADE less than 0.02 and 0.01, respectively.

	MADE < 0.05		MADE < 0.02		MADE < 0.01	
	CPU time	n_u	CPU time	n_u	CPU time	n_u
$p = 2$	0.00013	8	0.00174	12	0.00942	18
$p = 3$	0.00023	12	0.00429	27	0.04488	40
$p = 4$	0.00023	16	0.00958	81	0.11631	122
$p = 5$	0.00043	20	0.02991	153	0.35146	345
$p = 6$	0.00054	24	0.07419	274	0.90636	1386
$p = 8$	0.00334	72	0.34695	1228	9.83845	9324
$p = 10$	0.01361	90	1.54246	3450	104.45206	88412

The results are shown in Table 2. More detailed results are given in Figures 14 and 15 in supplementary material S.2. By comparing Tables 1 and 2, we see that the number of directional vectors and CPU time needed increase when the variables of X are dependent.

Let X still represent the multivariate normally distributed variable with covariances (10). Table 3 shows results for the multivariate lognormal distribution $Y = \exp(X)$. More detailed results are given in Figures 16 and 17 in supplementary material S.2.

Tables 2 and 3 show that a specific level of MADE is reached faster for the lognormal distribution than for the multivariate distribution documenting that the procedure efficiently can characterize non-Gaussian distributions.

To the best of our knowledge, the algorithm by [23] is the most efficient algorithm in the literature to estimate Tukey α -depth regions. The authors focus on estimating exact trimmed α -depth regions resulting in complex combinatorial algorithms and the computation burden explodes with the number of samples. In comparison, the computational complexity of our algorithm increases linearly with the number of samples. The authors can document estimation results up to dimension $p = 9$, but only when the number of samples are restricted to less than 80. The algorithm by [23] is not constructed to handle streaming data.

5.2. Synthetic experiments - Offline setting

In this section, we compare the performance of the incremental quantile estimator, DUMIQE, with state-of-the-art offline quantile estimators to estimate α -depth regions when data is known in advance. State-of-the-art offline quantile estimators are based on using weighted averages of consecutive order statistics

$$Q(\alpha) = (1 - \delta)y[j] + \delta y[j + 1]$$

where $\frac{j-m}{N} \leq \alpha < \frac{j-m+1}{N}$, $y[j]$ is the j th order statistic of the sample, m a constant and N the sample size. We use $m = \frac{\alpha+1}{3}$ and $\delta = N\alpha + m - j$ and define $\alpha[k] = \frac{k-1/3}{N+1/3}$. The sample quantiles can be read from a linear interpolation between the points $(\alpha[k], y[k]), k = 1, \dots, N$. The resulting quantile estimates are approximately median-unbiased regardless of the distribution of the data. This is the method referred to as Type 8 in the quantile function in R and is the one recommended by [16].

We consider the multivariate normal distribution case with covariance matrix as given in (10), sample sizes $N = 500, 2000, 10^4$ and $5 \cdot 10^4$ and dimensions $p = 2$ and $p = 3$. For $p = 2$ and $p = 3$, we used 1500 and 7500 directional vectors, respectively, which were sufficiently many to obtain satisfactory performance.

The results are shown in Table 4. We see that the estimation errors using DUMIQE are about 1.5 time that of the offline estimator. If fewer directional vectors were used, the differences in estimation error were substantially reduced. Further, the computational time of the offline estimator is about ten times larger than the DUMIQE

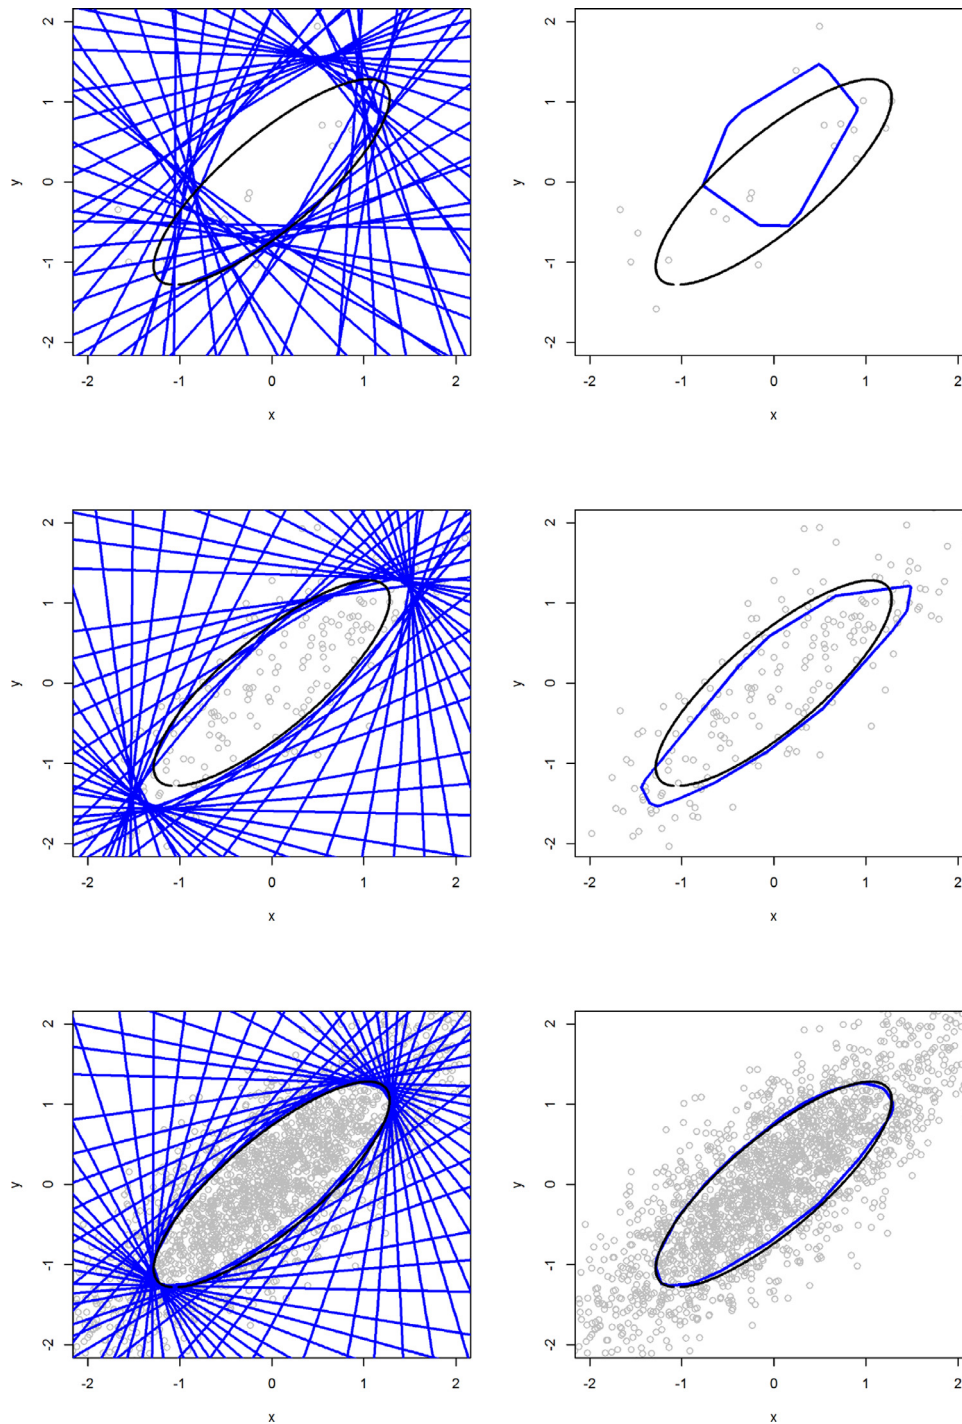


Fig. 3. Multivariate normal distribution case. Estimation of α -depth region for $\alpha = 0.1$ using $n_d = 50$ directional vectors. The rows from top to bottom show the estimates for 20, 200 and 2000 observations. The left and right column show all the half planes and the resulting envelopes in blue, respectively. The black curves show the true α -depth contour. (For interpretation of the references to colour in this figure legend, the reader is referred to the web version of this article.)

estimator. In other words, if computational time or memory usage are not an issue, the offline estimator combined with a large amount of directional vectors will give the most precise estimates from the samples. Otherwise, incremental quantile estimators are preferable even for offline settings.

5.3. Synthetic experiments - Dynamically changing data streams

In this section, we consider the problem of tracking α -depth regions of dynamically varying data streams. Fig. 4 illustrates the

problem. In each panel, the expectation vector of the data stream moved from the bottom left to the upper right. At the same time the correlation, changed from strongly positive, 0.8, to strongly negative, -0.8 . For the 10^3 samples case (first row), the algorithm was able to track the α -depth regions satisfactory. With 10^4 samples, the estimates improve significantly and with 10^5 observations, the estimates are very close to the true contours. With 10^4 and 10^5 samples, 50 directional vectors give better and smoother estimates than 10 directional vectors.

Table 2

Multivariate normal distribution case: The second and third columns show the CPU time (in seconds) and the number of directional vectors used to obtain a mean absolute depth error (MADE) less than 0.05, respectively.

	MADE < 0.05		MADE < 0.02		MADE < 0.01	
	CPU time	n_u	CPU time	n_u	CPU time	n_u
$p = 2$	0.00034	18	0.00622	40	0.03734	40
$p = 3$	0.00095	27	0.03003	135	1.38903	135
$p = 4$	0.00238	54	0.13145	274	7.47343	616
$p = 5$	0.01275	102	0.43698	777	8.47334	3936
$p = 6$	0.03603	183	1.81652	3118	45.88106	15,786
$p = 8$	0.17285	819	23.21460	20,979	988.12085	358,438
$p = 10$	0.68053	2300	245.91893	198,927	-	-




Fig. 4. Tracking of α -depth contours for $\alpha = 0.05, 0.2$ and 0.4 : In each panel the gray dots are outcomes from the data stream. The first observations from the data stream are shown in dark gray and the dots become lighter gray as time progresses. The left and right column show cases with $n_u = 10$ and 50 directional vectors, respectively. The rows from top to bottom show cases with a total for $10^3, 10^4$ and 10^5 observations, respectively.

Evaluation for $p > 2$ is given below. Due to the computational burden of evaluating estimation error of non-elliptic distributions, the analysis was restricted to Gaussian distributions. Let $X_n = (X_{n,1}, \dots, X_{n,p})^T$ be multivariate normally distributed with

$$\mu_{n,i} = E(X_{n,i}) = \sin\left(\frac{2\pi}{T}n + \psi_i\right), \quad i = 1, \dots, p \quad (11)$$

where $\psi_i, i = 1, \dots, p$ are independent uniformly distributed variables on the interval $[0, 2\pi]$ ensuring that the marginal expectations are out of phase. Covariance between $X_{n,i}$ and $X_{n,j}$ is

$$\text{Cov}(X_{n,i}, X_{n,j}) = \left(0.4 \sin\left(\frac{2\pi}{T}n + \psi\right) + 0.4\right)^{|i-j|} \quad (12)$$

where ψ is uniformly distributed on the interval $[0, 2\pi]$.

Table 3

Multivariate lognormal distribution case: The second and third columns shows the CPU time (in seconds) and the number of directional vectors used to obtain a mean absolute depth error (MADE) less than 0.05, respectively. The fourth and fifth and the sixth and seventh columns show the same to obtain MADE less than 0.02 and 0.01, respectively.

	MADE < 0.05		MADE < 0.02		MADE < 0.01	
	CPU time	n_u	CPU time	n_u	CPU time	n_u
$p = 2$	0.00013	8	0.00957	27	0.11169	40
$p = 3$	0.00024	27	0.01533	135	0.56418	202
$p = 4$	0.00021	24	0.03214	274	1.64312	924
$p = 5$	0.00043	45	0.14592	1166	6.16044	3936
$p = 6$	0.00053	54	0.27431	2079	9.30407	15,786

Table 4

Offline experiment: Comparison of the DUMIQE estimator and the estimator recommended in [16] to estimate α -depth contours for $\alpha = 0.05, 0.2$ and 0.4 . MADE, ED and CPU refers to the error measures in (7) and (8) (multiplied by 10^3) and CPU time used (in seconds), respectively. N refers to the sample size.

N	Method	$p = 2$			$p = 3$		
		MADE	ED	CPU	MADE	ED	CPU
500	Offline	14.9	43.1	0.291	16.9	40.5	1.634
	DUMIQE	25.1	63.9	0.045	34.9	69.7	0.288
2000	Offline	7.0	20.7	1.421	7.2	18.2	9.489
	DUMIQE	10.6	28.5	0.182	12.2	26.5	1.154
10^4	Offline	3.0	9.0	8.761	3.0	7.7	55.12
	DUMIQE	4.4	12.1	0.908	4.6	10.6	5.769
$5 \cdot 10^4$	Offline	1.3	4.0	52.32	1.3	3.5	326.0
	DUMIQE	1.8	5.4	4.542	2.0	4.7	28.84

Tables 5 to 6 show results tracking α -depth regions for $\alpha = 0.05, 0.2$ and 0.4 for periods $T = 10^3$ and $T = 10^4$ under optimal choices of the tuning parameter¹ More detailed results are given in Figures 18 and 19 in supplementary material S.3. For $T = 10^3$, MADE is around 0.05 and the estimation error does not decrease with increasing number of directional vectors which may seem surprising. The reason is that if the quantile estimates are poor, the intersections of the resulting halfspaces do not necessarily become better by adding more halfspaces. For $T = 10^4$ MADE is between 0.02 and 0.03. The optimal number of halfspaces increases with dimension, but not dramatically.

The algorithm is computationally very efficient. For dimension $p = 5$ the algorithm can optimally process 10^4 to 10^5 observations from a data stream every second on a single CPU processor.

By using more equidistant directional vectors, we expect reduction in the tracking error. Consider the dynamic case above except that the directional vectors are chosen more equidistantly. Directional vectors were generated using the filtering procedure in Section 3 with $N_u = 10n_u$.

The results are shown in Table 7 and more detailed results are given in Figure 21 in supplementary material S.3.

By comparing Tables 5 and 6 with 7, we see that for $T = 10^3$ and $T = 10^4$, minimum MADE is reduced from 0.045 to 0.040 and from 0.0226 to 0.0216, respectively. However, more importantly, by using equidistant vectors, the best results are obtained using fewer directional vectors. For both $T = 10^3$ and $T = 10^4$, the optimal number of vectors are reduced from 25 to 10. Finally, we observe significant improvement if only five directional vectors were used. Using equidistant directional vectors adds an additional computational cost in the initialization of the algorithm, but results into gained peak performance and fewer directional vectors, and

¹ In a practical situation, the history of the data stream can be used to estimate (or track) optimal values of the tuning parameters. We are currently working on such procedures.

thus less computation time and memory are needed during tracking.

6. Real-life data examples

In this section, we use the algorithm on a real-life dataset related to activity change detection. A second real-life data example related to real-time event detection using Twitter data is given in supplementary material S.4.

We demonstrate how the algorithm can be used to detect outliers and events and perform classifications in dynamic settings. For example, related to event detection, by characterizing a data stream distribution with multiple depth contours, in practice any change in the data stream distribution can be detected. Not only changes in common properties such as expectation and covariance structure, but also changes in shape such as a change from an elliptic to a non-elliptic distribution.

6.1. Activity change detection

Activity recognition is a highly active field of research where sensory information is used to automatically detect and identify activities of users. Activity recognition can help for example detect sedentary lifestyle and prompt the user to perform healthy exercises.

We consider an accelerometer dataset from the WISDM (Wireless Sensor Data Mining) project [22]. Accelerations in x, y and z directions were observed, with a frequency of 20 observations per second, while users were performing the activities walking, jogging, walking up a stairway and walking down a stairway. A total of 36 users were observed and the dataset contains a total of 989875 observations.

Current research focuses on supervised approaches where historic and annotated activity observations are used to train an activity classification model. E.g.[22], trained models such as decision trees and neural networks. However, such an approach is highly sensitive to any temporal changes in the data, e.g. if the user switches to an activity that is not part of the training material as a consequence for example of becoming fitter, sick etc. In this example we rather take an unsupervised approach and the goal is to detect whenever the user changes activity. Since we receive 20 accelerometer observations per second, it is important that the streaming approach is computationally efficient.

Fig. 5 shows in gray x, y and z acceleration for an arbitrary user. The red lines show when the user changed activity. Acceleration distributions are fairly stationary within an activity, but with some gradual and abrupt changes. The users changed activities in many cases as often as every 30 seconds making this a challenging tracking and change detection problem. Fig. 6 shows scatter plots for two arbitrary sessions with minimal temporal trend. The simultaneous acceleration distributions vary substantially between sessions and are often far from being elliptical. Further, even though the distributions are different, the mean and covariances are often quite similar making the change detection task based on elliptic distributions (Mahalanobis distance) challenging. We thus suggest the following depth based change detection procedure:

1. Track α -depth contours of the simultaneous acceleration distribution by tracking n_u directional quantiles using the DUMIQE algorithm with tuning parameter λ .
2. Compute the Euclidean distance between the current α -depth contours and the contours h seconds back in time using Equation (8), denoted ED_t at time t .
3. Track the expectation and standard deviation of ED_t distribution using exponential moving average

$$E(ED_t) = (1 - \delta)E(ED_{t-1}) + \delta ED_t$$

Table 5

Tracking of α -depth regions for $\alpha = 0.05, 0.2$ and 0.4 for the distribution characterized by (11) and (12) with a period $T = 10^3$. The columns 'Freq' refer to how many times per millisecond the algorithm can update an α -depth region when running on a single 1.8 GHz CPU processor.

n_u	$p = 2$		$p = 3$		$p = 4$		$p = 5$	
	MADE	Freq	MADE	Freq	MADE	Freq	MADE	Freq
5	0.0559	972.7	-	-	-	-	-	-
10	0.0475	478.9	0.0577	486.7	-	-	-	-
25	0.0445	189.2	0.0457	189.7	0.0510	184.7	-	-
50	0.0474	95.2	0.0467	95.1	0.0492	93.3	0.0521	92.4
100	0.0504	47.9	0.0502	47.4	0.0514	46.8	0.0523	46.4
200	-	-	0.0536	23.4	0.0546	22.6	0.0541	22.7
500	-	-	-	-	0.0590	9.1	0.0576	8.9
1000	-	-	-	-	-	-	0.0604	4.5

Table 6

Tracking of α -depth regions for $\alpha = 0.05, 0.2$ and 0.4 for the distribution characterized by (11) and (12) with a period $T = 10^4$. The columns 'Freq' refer to how many times per millisecond the algorithm can update an α -depth region when running on a single 1.8 GHz CPU processor.

n_u	$p = 2$		$p = 3$		$p = 4$		$p = 5$	
	MADE	Freq	MADE	Freq	MADE	Freq	MADE	Freq
5	0.0439	976.3	-	-	-	-	-	-
10	0.0305	480.1	0.0499	484.3	-	-	-	-
25	0.0226	189.2	0.0318	188.5	0.0429	184.7	-	-
50	0.0227	95.4	0.0277	94.3	0.0342	93.6	0.0395	91.1
100	0.0236	48.0	0.0275	47.3	0.0306	47.0	0.0337	46.1
200	-	-	0.0289	23.3	0.0299	22.7	0.0312	22.6
500	-	-	-	-	0.0307	9.1	0.0309	9.0
1000	-	-	-	-	-	-	0.0316	4.5

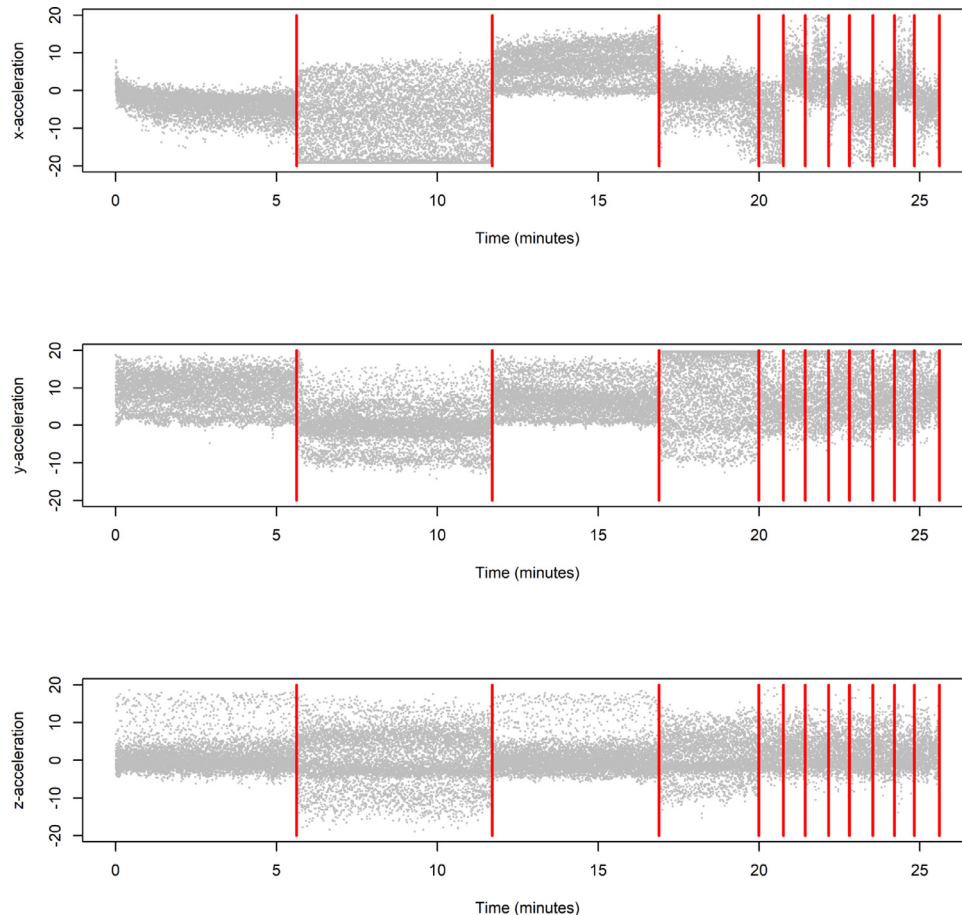


Fig. 5. The gray dots show accelerometer observations for an arbitrary user. The red lines show when the user changes activity. (For interpretation of the references to colour in this figure legend, the reader is referred to the web version of this article.)

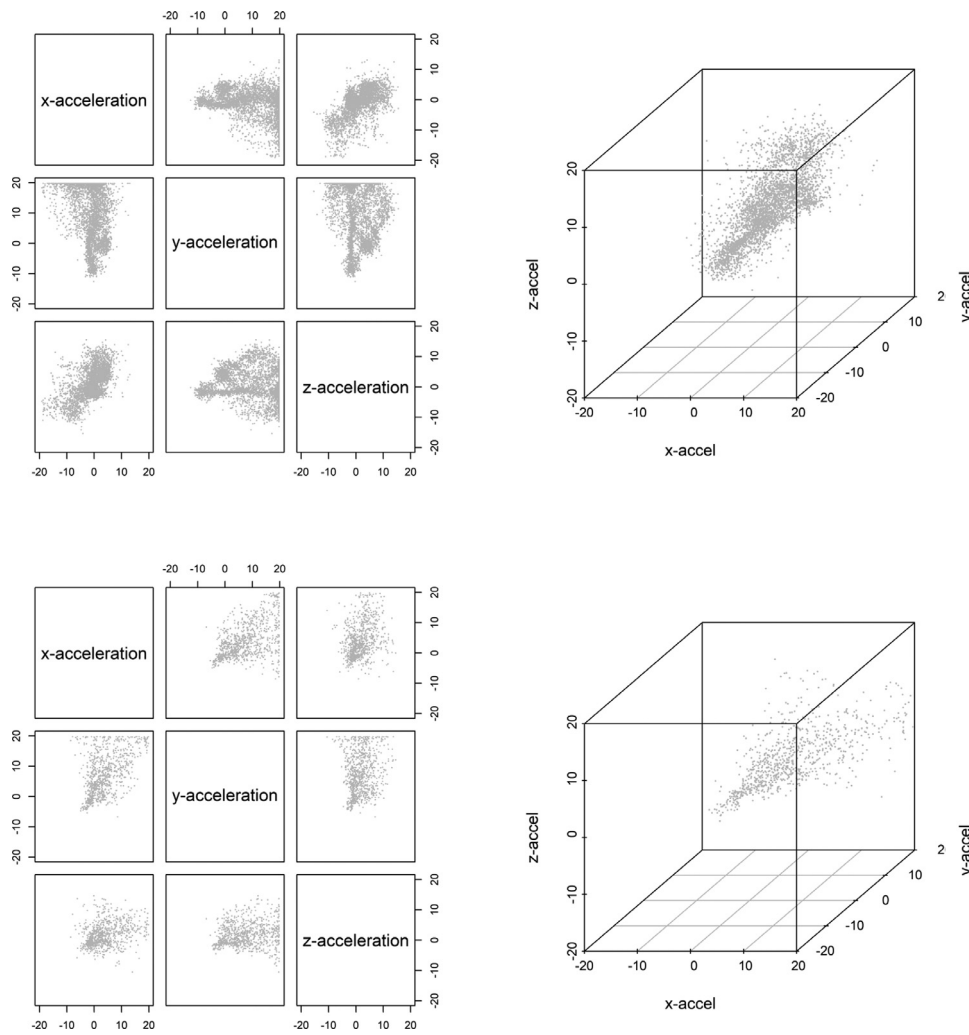


Fig. 6. The first and the second row show scatterplot of accelerometer observations for two activity sessions.

Table 7
Tracking of α -depth regions for $\alpha = 0.05, 0.2$ and 0.4 for the distribution characterized by (11) and (12) using fairly equidistant directional vectors. Tracking error is measured using MADE. The left and right columns show results for $T = 10^3$ and 10^4 , respectively. Dimension is $p = 2$.

n_u	$T = 10^3$	$T = 10^4$
5	0.0465	0.0303
10	0.0400	0.0216
25	0.0440	0.0217
50	0.0477	0.0226
100	0.0505	0.0236

$$E(ED_t^2) = (1 - \delta)E(ED_{t-1}^2) + \delta ED_t^2$$

$$SD(ED_t) = \sqrt{E(ED_t)^2 - E(ED_t^2)}$$

- When the user changes activity, we expect ED_t to rapidly increase. A new activity is detected when ED_t is more than η standard deviations higher than $E(ED_t)$, i.e. $ED_t \geq E(ED_t) + \eta SD(ED_t)$.

- When a new activity is detected, restart the tracking of the α -depth contours and go back to step 1.

This approach is elegant since by virtue of measuring difference in depth contours, it can detect virtually any kind of change in the shape of the simultaneous acceleration distribution, for example a change from an elliptic to a non-elliptic distribution. Given the properties of the observations in this application, this flexibility is important.

We compare the approach against an identical approach except that in the first part of the algorithm the mean and covariance structure (and not depth contours) were tracked using multivariate exponentially weighted moving average (MEWMA) [24].

We measured the performance of the depth and the MEWMA approaches for a wide range of values for the tuning parameters. As several sessions lasted for only 30 seconds, it was thus important for the tracking algorithms to rapidly adapt to a session before a new change of activity took place. In the first step of the procedures, we thus chose decreasing values of the tuning parameters, but with a minimum value to take into account the dynamic changes in accelerations within a session, $\lambda_t = \max\{1/t, \lambda_{\min}\}$, and tried the values 0.1, 0.05 and 0.01 for λ_{\min} ². This performed better than using constant values of the tuning parameter. We further

² For MEWMA, λ_t refers to the moving average tuning parameter and $\lambda_t = 1/t$ is thus equivalent to the sample mean.

Table 8
Change detection example. Results for the depth approach.

λ_{\min}	δ_{\min}	h	η	n_u	Precision	Recall	F1 score	Det. delay (sec)
0.01	0.01	100	8	20	0.796	0.497	0.612	1.325
0.01	0.01	100	8	50	0.781	0.486	0.599	1.150
0.01	0.05	200	8	20	0.532	0.682	0.597	1.415
0.01	0.05	50	8	50	0.613	0.564	0.587	1.876
0.01	0.01	200	8	20	0.735	0.488	0.587	1.202
0.01	0.05	200	8	50	0.510	0.673	0.580	1.482
0.01	0.01	200	8	50	0.719	0.480	0.575	1.219
0.01	0.05	100	8	20	0.538	0.618	0.575	1.010
0.05	0.10	200	8	20	0.539	0.616	0.575	1.903
0.01	0.05	50	8	20	0.613	0.532	0.570	1.894

Table 9
Change detection example. Results for the MEWMA approach.

λ_{\min}	δ_{\min}	h	η	Precision	Recall	F1 score	Det. delay (sec)
0.01	0.01	200	8	0.454	0.697	0.550	1.553
0.05	0.01	200	8	0.447	0.697	0.545	1.691
0.05	0.01	50	8	0.438	0.699	0.539	2.249
0.01	0.01	50	8	0.421	0.711	0.529	1.736
0.01	0.01	100	8	0.398	0.737	0.517	1.293
0.05	0.01	100	8	0.388	0.711	0.502	1.747
0.05	0.05	200	8	0.353	0.760	0.483	1.525
0.05	0.05	50	8	0.336	0.818	0.476	1.832
0.05	0.10	200	8	0.336	0.803	0.474	1.522
0.01	0.05	200	8	0.330	0.777	0.463	1.281

tried the values 0.1, 0.05 and 0.01 for δ , 2.5, 5 and 10 seconds for h and 2, 5 and 8 for η . Further, for the depth approach we used three depth contours with α equal to 0.2, 0.05 and 0.01 and tried $n_u = 20$ or 50 directional vectors. We ran the two change detection approaches for the whole dataset for all the combinations of the parameters. This resulted in a total of 162 and 81 experiments for the depth and the MEWMA approaches, respectively.

Precision, recall and the F1 score were used to measure performance [31]. If the approach detects more than one change between two true changes, we characterize the first change as a correct detection and the others as false detections and define

$$\text{Precision} = \frac{\text{No. of correct detections}}{\text{No. of detections}}$$

$$\text{Recall} = \frac{\text{No. of correct detections}}{\text{No. true changes}}$$

$$\text{F1 score} = \frac{2 \cdot \text{Precision} \cdot \text{Recall}}{\text{Precision} + \text{Recall}}$$

Tables 8 and 9 show the top ten results with respect to the F1 score. The depth approach outperforms the MEWMA with respect to the F1 score and in addition detects the true changes more rapidly. The performance of the depth approach does not seem to be particularly sensitive to the number of directional quantiles used.

7. Closing remarks

In this paper we have presented a computationally and memory efficient procedure to estimate and track Tukey α -depth contours using incremental quantile estimators. The algorithms use the results by [20] according to which the α -depth region equals the directional quantile envelope. We further demonstrated how incremental quantile estimators can be used to efficiently estimate the directional quantile envelope. By using incremental quantile estimators, we are able to recursively estimate and track α -depth contours, and to the best of our knowledge, it is the first method in the literature with this ability. However, as shown in Section 5.2,

if the amount of data is limited, it is better to estimate α -depth contours using traditional offline quantile estimators.

The algorithms estimated Tukey depth contours equally well for both elliptic (Gaussian) and non-elliptic distributions. The performance, however, depends on the degree of curvature for the true depth contours being closely related to the degree of dependency between variables. For static data streams, the algorithm estimated a depth contour of dimension $p = 10$ with a mean absolute error in Tukey depth less than 0.01 in 1.2 and 125 minutes for independent and strongly dependent variables, respectively, on a single CPU processor. For dynamically changing data streams, even for dimensions as high as $p = 5$, the algorithm was able to process tens of thousands of observations per second and track depth contours with high precision. We have not found any studies that have been able to estimate depth contours of such a high dimension and for such a large amount of data, which documents the efficiency of the algorithm.

The real-life data examples demonstrate that the procedure is useful to track and detect changes in complex distributional patterns.

To estimate α -depth contours, the number of directional vectors, n_u , and values of tuning parameters in the incremental quantile tracking algorithms must be chosen. We are currently working on procedures that use information from the history of the data stream to recursively update such values. Tukey depth is best suited to account for convex features of the distribution of interest. However, there exist other modified depth measures that better account for non-convex features [6]. In the future, we plan to extend the method in this paper in order to also be applied to these depth measures.[1,2,12,26]

Declaration of Competing Interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

Supplementary material

Supplementary material associated with this article can be found, in the online version, at doi:10.1016/j.patcog.2021.108339.

References

- [1] N. Alkhamees, M. Fasli, Event detection from social network streams using frequent pattern mining with dynamic support values, in: Big Data (Big Data), 2016 IEEE International Conference on, IEEE, 2016, pp. 1670–1679.
- [2] F. Atefeh, W. Khreich, A survey of techniques for event detection in twitter, Comput Intell 31 (1) (2015) 132–164.
- [3] J.O. Cerdeira, T. Monteiro-Henriques, M.J. Martins, P.C. Silva, D. Alagador, A.M. Franco, M.L. Campagnolo, P. Arsénio, F.C. Aguiar, M. Cabeza, Revisiting niche fundamentals with tukey depth, Methods Ecol. Evol. 9 (12) (2018) 2349–2361.
- [4] Y. Chang, Z. Tu, W. Xie, B. Luo, S. Zhang, H. Sui, J. Yuan, Video anomaly detection with spatio-temporal dissociation, Pattern Recognit (2021) 108213.

- [5] F. Chebana, T.B. Ouarda, Depth-based multivariate descriptive statistics with hydrological applications, *Journal of Geophysical Research: Atmospheres* 116 (D10) (2011).
- [6] V. Chernozhukov, A. Galichon, M. Hallin, M. Henry, et al., Monge–kantorovich depth, quantiles, ranks and signs, *Ann Stat* 45 (1) (2017) 223–256.
- [7] D. Eddelbuettel, *Seamless rand c++ integration with rcpp*, Springer, New York, 2013. ISBN 978-1-4614-6867-7
- [8] D. Eddelbuettel, R. François, *Rcpp: seamless r and c++ integration*, *J Stat Softw* 40 (8) (2011) 1–18.
- [9] S.M. Erfani, S. Rajasegarar, S. Karunasekera, C. Leckie, High-dimensional and large-scale anomaly detection using a linear one-class svm with deep learning, *Pattern Recognit* 58 (2016) 121–134.
- [10] H.L. Hammer, A. Yazidi, H. Rue, A new quantile tracking algorithm using a generalized exponentially weighted average of observations, *Applied Intelligence* 49 (4) (2019) 1406–1420.
- [11] H.L. Hammer, A. Yazidi, H. Rue, Joint tracking of multiple quantiles through conditional quantiles, *Inf Sci (Ny)* 563 (2021) 40–58.
- [12] M. Hasan, M.A. Orgun, R. Schwitter, A survey on real-time event detection from the twitter data stream, *Journal of Information Science* (2017). 0165551517698564
- [13] S. Huang, Z. Kang, Z. Xu, Q. Liu, Robust deep k-means: an effective and simple method for data clustering, *Pattern Recognit* 117 (2021) 107996.
- [14] M. Hubert, P. Rousseeuw, P. Segaeert, Multivariate and functional classification using depth and distance, *Adv Data Anal Classif* 11 (3) (2017) 445–466.
- [15] M. Hubert, P.J. Rousseeuw, P. Segaeert, Multivariate functional outlier detection, *Statistical Methods & Applications* 24 (2) (2015) 177–202.
- [16] R.J. Hyndman, Y. Fan, Sample quantiles in statistical packages, *Am Stat* 50 (4) (1996) 361–365.
- [17] B.K. Iwana, S. Uchida, Time series classification using local distance-based features in multi-modal fusion networks, *Pattern Recognit* 97 (2020) 107024.
- [18] R. Jörnsten, Clustering and classification based on the l1 data depth, *J Multivar Anal* 90 (1) (2004) 67–89.
- [19] S. Kim, B.M. Mun, S.J. Bae, Data depth based support vector machines for predicting corporate bankruptcy, *Applied Intelligence* 48 (3) (2018) 791–804.
- [20] L. Kong, I. Mizera, Quantile tomography: using quantiles with multivariate data, *Stat Sin* (2012) 1589–1610.
- [21] D. Kosiorowski, Z. Zawadzki, *Depthproc an r package for robust exploration of multidimensional economic phenomena*, arXiv preprint arXiv:1408.4542 (2014).
- [22] J.R. Kwapisz, G.M. Weiss, S.A. Moore, Activity recognition using cell phone accelerometers, *ACM SigKDD Explorations Newsletter* 12 (2) (2011) 74–82.
- [23] X. Liu, K. Mosler, P. Mozharovskyi, Fast computation of tukey trimmed regions and median in dimension $p > 2$, *Journal of Computational and Graphical Statistics* (2019) 1–31.
- [24] C.A. Lowry, W.H. Woodall, C.W. Champ, S.E. Rigdon, A multivariate exponentially weighted moving average control chart, *Technometrics* 34 (1) (1992) 46–53.
- [25] J. Ma, Y. Zhang, L. Zhang, Discriminative subspace matrix factorization for multiview data clustering, *Pattern Recognit* 111 (2021) 107676.
- [26] J.-C. Massé, Asymptotics for the tukey depth process, with an application to a multivariate trimmed mean, *Bernoulli* (2004) 397–419.
- [27] K. Mosler, *Depth Statistics*, in: *Robustness and complex data structures*, Springer, 2013, pp. 17–34.
- [28] R Core Team, *R: A Language and Environment for Statistical Computing*, R Foundation for Statistical Computing, 2021. Vienna, Austria
- [29] N. Rastin, M.Z. Jahromi, M. Taheri, A generalized weighted distance k-nearest neighbor for multi-label problems, *Pattern Recognit* 114 (2021) 107526.
- [30] E.B. Saff, A.B. Kuijlaars, Distributing many points on a sphere, *The mathematical intelligencer* 19 (1) (1997) 5–11.
- [31] M. Sokolova, G. Lapalme, A systematic analysis of performance measures for classification tasks, *Information Processing & Management* 45 (4) (2009) 427–437.
- [32] J.W. Tukey, Mathematics and the picturing of data, in: *Proceedings of the international congress of mathematicians*, 2, 1975, pp. 523–531.
- [33] B. Williams, M. Toussaint, A.J. Storkey, Modelling motion primitives and their timing in biologically executed movements, in: *Advances in neural information processing systems*, 2008, pp. 1609–1616.
- [34] A. Yazidi, H. Hammer, Multiplicative update methods for incremental quantile estimation, *IEEE Trans Cybern* 49 (3) (2017) 746–756.
- [35] V. Zavrtnik, M. Kristan, D. Skočaj, Reconstruction by inpainting for visual anomaly detection, *Pattern Recognit* 112 (2021) 107706.
- [36] Y. Zuo, R. Serfling, General notions of statistical depth function, *Ann Stat* (2000) 461–482.



Hugo Lewi Hammer received the M.Sc. and Ph.D. degrees from the Norwegian University of Science and Technology, in 2003 and 2008, respectively. He is currently a professor with the Department of Computer Science, OsloMet Oslo Metropolitan University, Oslo, Norway and an adjunct research scientist at Simula Metropolitan Centre, Oslo, Norway. Before joining OsloMet and SimulaMet, he worked as a researcher with Norwegian Computer Center, Oslo, Norway. His research interests include computer intensive statistical methods, machine learning, learning automata and stochastic optimization.



Anis Yazidi received the M.Sc. and Ph.D. degrees from the University of Agder, Grimstad, Norway, in 2008 and 2012, respectively. He is currently a professor with the Department of Computer Science OsloMet Oslo Metropolitan University, Oslo, Norway. Before joining OsloMet, he worked as a researcher with Teknova AS, Grimstad, Norway. His current research interests include machine learning, learning automata, stochastic optimization, software-defined networks and cloud computing. He led the research group Autonomous Systems and Networks at OsloMet from 2015 to 2018. He is currently the leader of the research group Applied Artificial Intelligence at OsloMet.



Håvard Rue received the M.Sc. and Ph.D. degrees from the Norwegian Institute of Technology, in 1988 and 1993, respectively. He is currently a professor in statistics with the Computer, Electrical and Mathematical Science and Engineering Division, King Abdullah University of Science and Technology (KAUST), Saudi Arabia. Before joining KAUST, he worked as a professor in statistics at Norwegian University of Science and Technology. His research interests lie in computational Bayesian statistics and Bayesian methodology such as priors, sensitivity and robustness. His main body of research is built around the R-INLA project (www.r-inla.org), which aims to provide a practical tool for approximate Bayesian analysis of latent Gaussian models, often at extreme data scales. This project also includes efforts to use stochastic partial differential equations to represent Gaussian fields, for the use in spatial statistics. He is the leader of the research group Bayesian Computational Statistics and Modeling.