

# GUC: Unsupervised non-parametric Global Clustering and Anomaly Detection

**Chris Solomou**

CS2291@YORK.AC.UK

*University of York*

*Deramore Lane Heslington York*

*York, UK*

## Abstract

Clustering is a crucial task in the fields of Optimization and Machine Learning, with the ability of grouping similar data points together being of particular importance. In this paper, we propose a novel non-parametric algorithm that performs Global Clustering and Anomaly Detection in an unsupervised manner. Our algorithm is both effective and efficient, requiring no prior assumptions or domain knowledge to be applied. It features two modes that utilize the distance from the dataset’s center for clustering data points together. The first mode splits the dataset into global clusters where each cluster signifies proximity from the center. The second mode employs a threshold value for splitting the points into outliers and inliers. We evaluate our proposal against other prominent methods using synthetic and real datasets. Our experiments demonstrate that the proposed algorithm achieves state-of-the-art performance with minimum computational cost.

## 1. Introduction

We begin our work by defining a vector  $x = (x_1, x_2, \dots, x_d)$  to be a data point in the  $d$ -dimensional Euclidean space  $x \in \mathbb{R}^d$  drawn *iid* from an unknown distribution  $X$ . Here, the pdf of the distribution is also unknown, hence the probability of assuming a particular point. In the context of outlier detection, a supervised clustering algorithm such as KNN would generally perform well for a set of data  $X = \{(x_1, y_1), (x_2, y_2), \dots, (x_N, y_N)\}$  with known labels  $y \in \{0, 1\}$ , due to its ability to predict labels based on neighboring points. However, in unsupervised clustering the labels are now latent with  $X$  being defined as  $X = \{(x_1, \bullet), (x_2, \bullet), \dots, (x_N, \bullet)\}$ , and splitting the data globally means that each point belongs to a cluster that is yet to be observed. Therefore, each point needs to be assigned to a particular cluster with center  $c \in C$  where  $C = \{c_1, c_2, \dots, c_K\}$  and  $K$  being a fixed quantity defined by the user. In our work we assume a non-binary classification of anomalies, and we aim to generate the set  $C$  consisting of  $K$  global centers, and assign each point to the nearest cluster. Global clusters and the distance from their centers are defined in the next sections. For each point  $x \in X$  the probability mass function of belonging to a cluster from  $\{c_1, c_2, \dots, c_K\}$  is defined by  $\sum_{k=1}^K P(c_k|x) = 1$  where  $P(c_k|x)$  needs to be estimated. From a probabilistic standpoint is difficult to estimate:

$$P(c_k|x) = \frac{P(x|c_k)P(c_k)}{P(c_1)P(x|c_1) + \dots + P(c_K)P(x|c_K)} = \frac{P(x|c_k)P(c_k)}{\sum_{k=1}^K P(c_k)P(x|c_k)}$$

since a distributional assumption would have to be made, and even so, the probability of observing a particular cluster is still unknown. We can however approximate  $P(c_k|x)$  with a function that assigns the point to the closest cluster-center, i.e.:  $P(c_k|x) := f(x|C) = \arg \min_{c \in C} \{d(x, c)\}$  with  $d(\cdot)$  representing a distance measure.

## 2. Background

Gaussian Mixture Models (GMM), leverage a set of data points  $\{x_i\}_{i=1}^N$ , assuming they arise from a mixture of  $K$  Gaussian components. When the model parameters  $\theta = (\pi_k, \{\mu_k, \sum_k\})$  for each cluster are known, the most probable cluster for a data point  $x$  can be estimated by  $c = \operatorname{argmax}_k P(c_k|x, \theta) \propto \operatorname{argmax}_k P(c_k|\theta)P(x|c_k, \theta)$ , using the Expectation Maximization (EM) algorithm. K-means [MacQueen et al. \(1967\)](#) is an iterative clustering algorithm that partitions a dataset, aiming to assign a point  $x$  to the closest cluster center, i.e.:  $c = \operatorname{argmin}_k \|x - \mu_k\|^2$ , where  $\mu_k$  represents the cluster center, given by the mean  $\mu_k = \frac{1}{N_{c_k}} \sum_{i=1}^{N_{c_k}} x_i$  of within-cluster points. Unsupervised, density-based clustering algorithms are designed to identify dense regions in data. DBSCAN - [Ester et al. \(1996\)](#) is governed by two hyperparameters epsilon ( $\epsilon$ ) and *MinPts*, and is able to form a density if the conditions of Maximality ( $\forall x, x' : \text{if } x \in c \text{ and } x' \text{ is density reachable from } x \Rightarrow x' \in c$ ) and Connectivity ( $\forall x, x' \in c : x \text{ is density-connected to } x'$ ) hold true w.r.t. to  $\epsilon$  and *MinPts* values. LOF - [Breunig et al. \(2000\)](#) improves density-based clustering by quantifying the degree that each point is isolated from the surrounding densities, s.t.  $LRD(x) = 1 / (\frac{\sum_{x' \in N_k(x)} d_k(x, x')}{|N_k(x)|})$ . LOF introduced the Outlier Factor score,  $LOF(x) = \sum_{x' \in N_k(x)} \frac{LRD_k(x')}{LRD_k(x)} / |N_k(x)|$ , whereas *LRD* denotes the local reachability density, and  $N_k(x)$  refers to the  $k$ -nearest neighbors of  $x$ . The One-Class SVM [Schölkopf et al. \(1999\)](#) is a binary classification algorithm that generates a decision boundary around the majority of the data, identifying outliers as deviations from this boundary. Mathematically, oSVM aims to solve the optimization problem:  $\operatorname{argmin}_{w, b, \xi} \frac{1}{2} \|w\|^2 + \frac{1}{vN} \sum_{i=1}^N \xi_i - b$  subject to the constraints  $w^T \phi(x_i) \geq b - \xi_i$  and  $\xi_i \geq 0$ , where  $\phi(\cdot)$  represents the RBF kernel and  $\xi_i$  a slack variable for margin formulation. Isolation Forests - [Liu et al. \(2008\)](#) leverage binary tree structures to isolate anomalies by recursively partitioning the data space. They employ an anomaly score  $S(x) = 2 \frac{E(h(x))}{c(N)}$  where  $E(h(x))$  is the expected path of  $x$  in a tree, and  $c(N)$  a normalization constant.

### 2.1. Definitions

Existing distance-based algorithms assign points to local clusters based on the distance from their centroids. In the case of global clustering, the center needs to be in a single fixed location, for measuring the topological distance of each point in the data space.

**Definition 1 - Global Center:** The center of a dataset is given by the set of average values of each dimension:  $\bar{x} = \{\bar{x}_j\}_{j=1}^d$ .

**Definition 2 - Global Clustering:** Given a set of global clusters centers  $C = \{c_1, c_2, \dots, c_K\}$  a point can be assigned to a cluster based on the minimum distance from the clusters' center, i.e.:  $c_x = \operatorname{argmin}_{c \in C} \{d(x, c)\}$ .

**Definition 3 - Global Outliers:** A point  $x$  is classified as an inlier if  $\|x - \bar{x}\|_2 \leq q$  where  $q$  is the threshold distance for a point to be considered an inlier.

## 3. Global Unsupervised (GU)-Clustering

In the context of Global Clustering, a point  $x$  can be topologically located from the dataset's center with some finite distance. Here we are interested in generating synthetic labels  $y \in \mathbb{R} \forall x \in X$  which indicate the proximity of each point from the center. Specifically using (3.1) as a distance measure,

$$y_i = \|x_i - \bar{x}\|_2 = \sqrt{\sum_{j=1}^d (x_i^{(j)} - \bar{x}_j)^2} \quad (3.1)$$

we measure the global distance of each point from the center. This process generates the vector  $Y$ , a univariate distribution of point-distances. This distribution results from the fitting step of our proposal, as outlined in Algorithm 1.

---

**Algorithm 1:** Point-Distance from Global Center

---

**Data:** Input:  $X \in \mathbb{R}^{N \times d}$

Initialize distance vector  $Y = (0, 0, \dots, 0) \forall x \in X$

**for**  $j \leftarrow 1$  **to**  $d$  **do**  
 |  $Y = Y + (X_j - \bar{X}_j)^2$   
**end**  
 $Y \leftarrow Y^{0.5}$

---

Furthermore, by utilizing the quantile discrepancy we can empirically measure the variability between a data point and a certain quantile. The vector of distances  $Y$  from (3.1), results in a univariate random variable defined on the real line  $\mathbb{R}$  with a strictly monotonically increasing CDF such that  $q(\theta) = F_Y^{-1}(\theta) = \inf\{y : F_Y(y) \geq \theta\}$  where  $\theta \in [0, 1]$  is a percentile and  $q(\theta)$  its corresponding quantile.

### 3.1. Outlier Detection

For performing outlier detection we can use the quantile discrepancy, and determine the value of  $q(\theta)$  that minimizes the following variability measure:

$$\theta \int_{y > q(\theta)} |y - q(\theta)| \frac{d}{d_y} F_Y(y) + (1 - \theta) \int_{y < q(\theta)} |y - q(1 - \theta)| \frac{d}{d_y} F_Y(y) \quad (3.1.1)$$

Equation (3.1.1) can be evaluated empirically, and provides a practical measure for quantifying the spread of the distribution at certain quantiles, by:

$$var(Y) = \frac{\sum_{i=1}^N I(y_i > q(\theta))}{\sum_{i=1}^N I(y_i)} \sum_{i=1}^N \theta |y_i - q(\theta)| + \frac{\sum_{i=1}^N I(y_i \leq q(\theta))}{\sum_{i=1}^N I(y_i)} \sum_{i=1}^N (1 - \theta) |y_i - q(1 - \theta)| \quad (3.1.2)$$

where  $\theta$  denotes the weight assigned to each quantile. Optimizing equation (3.1.2), we can determine the value of  $\theta$  that leads to the best separation between inliers and outliers based on the corresponding quantile. However, it may not be feasible to solve equation (3.1.2) analytically, and numerical optimization methods such as projected gradient descent are necessary (since  $\theta$  is constrained within the range  $[0, 1]$ ).

The quantile discrepancy for a single point can be mapped as  $f : \mathbb{R}^2 \rightarrow \{0, 1\}$ , where  $f(y, q(\theta)) = I(y > q(\theta))$  and for the vector of distances  $Y$ :

$$Y = f(Y, q(\theta)) := I(Y > q(\theta)) \quad (3.1.3)$$

where  $q(\theta)$  denotes the threshold distance for a point to be considered an inlier. Thus, we generate binary labels  $y \in \{0, 1\}$  to detect outliers based on their distance from the

global center. Points that exceed a certain threshold distance are classified as outliers, while the remaining points are considered inliers.

---

**Algorithm 2:** Global Outlier Detection

---

**Data:** Vector of distances  $Y = (y_1, y_2, \dots, y_N)$ ,  $q$   
 $Y \leftarrow I(Y > q)$

---

### 3.2. Global Clustering

For global clustering the same approach is adjusted for measuring the variability between a data point and  $K$  quantiles. Particularly, we aim to minimize the variability measure based on the number of chosen clusters  $K$ . Equation (3.1.1) can now be expressed as:

$$\sum_{k=1}^{K-1} \left(1 - \sum_{\substack{j=1 \\ j \neq k}}^K \theta_j\right) \int_{y \geq q(\theta_k)}^{q(\theta_{k+1})} |y - q(\theta_k)| \frac{d}{dy} F_Y(y) dy \quad (3.2.1)$$

Evaluating the integral, the overall variability of  $Y$  can be measured by:

$$\text{var}(Y) = \sum_{k=1}^{K-1} \left(1 - \sum_{\substack{j=1 \\ j \neq k}}^K \theta_j\right) |q(\theta_{k+1}) - q(\theta_k)| P(q(\theta_k) \leq Y \leq q(\theta_{k+1})) \quad (3.2.2)$$

Equation (3.2.2) can be optimized to determine the optimal number of clusters that best split a dataset. Future work aims to explore techniques such as reducing in-cluster variability (similar to the elbow method in k-means), or using gradient descent for finding the optimal value of  $K$ . For measuring the overall variability of a single point:

$$\sum_{k=1}^{K-1} \left(1 - \sum_{\substack{j=1 \\ j \neq k}}^K \theta_j\right) |q(\theta_{k+1}) - y| \quad (3.2.3)$$

Following, we can estimate  $K$  cluster centers by splitting  $Y$  into  $K$  quantiles, where the value of each quantile represents the center of each cluster.

$$C = \{c_1, c_2, \dots, c_K\} := \{q(\theta_1), q(\theta_2), \dots, q(\theta_K)\}$$

and for measuring the distance from each quantile (cluster-center):

$$\left\{ \left(1 - \sum_{\substack{j=1 \\ j \neq k}}^K \theta_j\right) |c_k - y| \right\} \quad \forall c_k \in C \quad (3.2.4)$$

Assuming a choice of equally spaced percentiles, we can omit the weights of each quantile, and the cluster for each point in  $Y$  can be determined based on the shortest distance from each cluster-center, such that:

$$c_x = \arg \min_{k \in \{1, \dots, K\}} \{|c_k - y|\} \quad (3.2.5)$$

Thus, we can assign the points to the closest cluster by measuring the distance from each cluster-center. Measuring the absolute distance from the center of each cluster, we are now able to group points from both ends of the data space, and permit them to be in the same cluster. A practical implementation of Global Clustering is presented in Algorithm 3.

---

**Algorithm 3:** Assigning Points to  $K$  Global Clusters

---

**Data:** Vector of distances  $Y = (y_1, y_2, \dots, y_N)$ ,  $K$

Calculate  $K$  quantiles from  $Y$

$Y \leftarrow \operatorname{argmin}_k \operatorname{abs}((q_1, q_2, \dots, q_K) - Y)$

---

## 4. Experiments and Results

### 4.1. Global Clustering in Synthetic Datasets

The proposed notion of Global Clustering is applied to low-dimensional synthetic datasets for visual inspection. Common metrics used to evaluate the quality of an unsupervised clustering algorithms, such as the silhouette coefficient and Dunn’s index are not applicable in global clustering since they measure the distance from within-cluster points. Therefore we measure the average distance of the points of each cluster from the dataset’s center,

$$\left\{ \frac{1}{N_{c_1}} \sum_{i=1}^{N_{c_1}} y_i, \frac{1}{N_{c_2}} \sum_{i=1}^{N_{c_2}} y_i, \dots, \frac{1}{N_{c_K}} \sum_{i=1}^{N_{c_K}} y_i \right\} \quad (4.1)$$

where  $Y_i$  as defined by equation (3.1). This allows us to perform global clustering and quantify the rank of each point as an outlier.

#### 4.1.1. SYNTHETIC DATASET 1. SINGLE-MODE DENSITY WITH 2 DIMENSIONS

Cluster	K-Means	DSCAN	GMM	GU-Clustering
1	0.3289	0.3653	0.4896	0.0695
2	0.3167	0.1610	0.374	0.1911
3	0.3429	0.2614	0.2742	0.3261
4	0.3977	0.1733	0.2806	0.5341

Table 1: Average distance from the global center

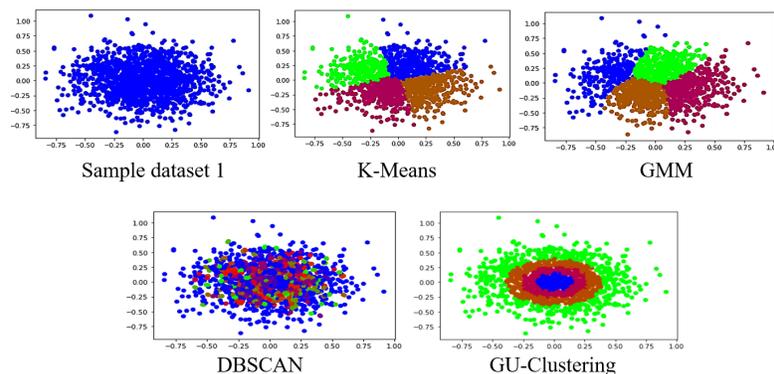


Figure 1: For the algorithms that allow cluster selection  $K$  was set to 4. DBSCAN generated a large number of clusters and does not produce meaningful results for either local or global clustering. GMM and K-Means split the data into 4 local clusters.

## 4.2. Outlier Detection

Although unsupervised anomaly detection does not utilize any label information, in this context they are needed for evaluation and comparison. For outlier detection, the rank of the anomalies should be considered. Specifically, each point should be ranked as being an outlier compared against the other data points that comprise the dataset. For evaluating predictive capabilities, the  $F_1$  score is employed due to the inherent class imbalance of outlier prediction datasets.

Dataset	Robust Covariance	oSVM	iForest	LOF	GU-Clustering
b-cancer	0.794	0.603	0.795	0.846	<b>0.988</b>
pen-global	0.601	0.660	0.750	0.706	<b>0.945</b>
letter	0.557	0.544	0.509	<b>0.701</b>	0.469
speech	0.5	0.505	0.508	<b>0.517</b>	0.492
sattelite	0.756	0.679	0.783	0.601	<b>0.993</b>
pen-local	0.5	0.5	0.5	0.5	0.5
annthyroid	<b>0.672</b>	0.502	0.527	0.492	0.482
shuttle	0.846	0.731	<b>0.978</b>	0.516	0.927
aloi	0.498	0.517	0.5	0.563	<b>0.605</b>
kdd99	0.502	0.665	0.761	0.501	<b>0.839</b>

Table 2: Comparison of the discriminative performance of each outlier detection algorithm. GU-Clustering performs well, yielding the highest AUC score (higher is better) in 5 of the 10 datasets that were used in this comparison.

Dataset	Robust Covariance	oSVM	iForest	LOF	GU-Clustering
b-cancer	0.978	0.972	0.989	<b>0.992</b>	0.988
pen-global	0.842	0.924	<b>0.944</b>	0.935	0.942
letter	0.896	0.942	0.939	0.963	<b>0.967</b>
speech	0.967	0.976	0.984	0.984	<b>0.991</b>
sattelite	0.986	0.991	<b>0.994</b>	0.988	0.993
pen-local	0.997	0.997	<b>0.999</b>	<b>0.999</b>	<b>0.999</b>
annthyroid	0.954	0.964	0.966	0.963	<b>0.981</b>
shuttle	0.989	0.990	<b>0.999</b>	0.982	0.991
aloi	0.941	0.971	0.970	0.974	<b>0.985</b>
kdd99	0.997	<b>0.999</b>	<b>0.999</b>	0.998	<b>0.999</b>

Table 3: Comparison of the predictive performance of each outlier detection algorithm. In terms of  $F_1$  score (higher is better) GU-clustering performs the best by either leading the results or being a close second.

The most prominent algorithms for outlier detection are compared against our proposal for measuring their performance. All of their hyperparameters were left at default with the contamination parameter set as the outlier percentage from Table 6 for achieving the best possible performance. For GU-Clustering, the hyperparameter  $q$  was set to 0.999. The value at  $q$  signifies the threshold distance for points to be classified as inliers. Preprocessing included applying a min-max transformation so the scale of all dimensions was between  $[0,1]$ . No other operations were performed with the datasets as presented by Goldstein and Uchida (2016), with further details provided in **Appendix A**.

## 5. Conclusions

Clustering is a critical process that aims to uncover the underlying structure of a set of data, which in turn can be harnessed to improve various aspects of optimization in Machine Learning. Despite its importance, past work in this area has not fully explored the potential of global clustering. In turn, we addresses the limitations of the current solutions by providing a novel and effective framework for both global clustering and anomaly detection. Using quantiles to split the distribution into  $K$  regions, each point is assigned to the value of the closest quantile (defined as the cluster center in the global clustering mode) or classified as an outlier if it exceeds the cutoff distance (outlier detection mode). Our proposal achieves state-of-the-art performance, is computationally efficient, has no underlying assumptions and can be applied to a wide range of datasets.

## 6. Acknowledgments

We are grateful and would like to thank our colleagues, Felix-Ulrich Oltean and Trinh Dong for their helpful comments and remarks.

## References

- Markus M Breunig, Hans-Peter Kriegel, Raymond T Ng, and Jörg Sander. Lof: identifying density-based local outliers. In *Proceedings of the 2000 ACM SIGMOD international conference on Management of data*, pages 93–104, 2000.
- Niko Brümmer, Sandro Cumani, Ondřej Glembek, Martin Karafiát, Pavel Matějka, Jan Pešán, Oldřich Plchot, Mehdi Souffar, Edward de Villiers, and Jan” Honza” Černocký. Description and analysis of the brno276 system for Ire2011. In *Odyssey 2012-the speaker and language recognition workshop*, 2012.
- Martin Ester, Hans-Peter Kriegel, Jörg Sander, Xiaowei Xu, et al. A density-based algorithm for discovering clusters in large spatial databases with noise. In *kdd*, volume 96, pages 226–231, 1996.
- Jan-Mark Geusebroek, Gertjan J Burghouts, and Arnold WM Smeulders. The amsterdam library of object images. *International Journal of Computer Vision*, 61(1):103–112, 2005.
- Markus Goldstein and Seiichi Uchida. A comparative evaluation of unsupervised anomaly detection algorithms for multivariate data. *PloS one*, 11(4):e0152173, 2016.
- Moshe Lichman et al. Uci machine learning repository, 2013.
- Fei Tony Liu, Kai Ming Ting, and Zhi-Hua Zhou. Isolation forest. In *2008 eighth ieee international conference on data mining*, pages 413–422. IEEE, 2008.
- James MacQueen et al. Some methods for classification and analysis of multivariate observations. In *Proceedings of the fifth Berkeley symposium on mathematical statistics and probability*, volume 1, pages 281–297. Oakland, CA, USA, 1967.
- Fabian Pedregosa, Gaël Varoquaux, Alexandre Gramfort, Vincent Michel, Bertrand Thirion, Olivier Grisel, Mathieu Blondel, Peter Prettenhofer, Ron Weiss, Vincent Dubourg, et al. Scikit-learn: Machine learning in python. *the Journal of machine Learning research*, 12:2825–2830, 2011.
- Bernhard Schölkopf, Robert C Williamson, Alex Smola, John Shawe-Taylor, and John Platt. Support vector method for novelty detection. *Advances in neural information processing systems*, 12, 1999.

## Appendix A. Datasets

The following datasets as presented by Goldstein et.al. [Goldstein and Uchida \(2016\)](#), are used for comparative evaluation of unsupervised anomaly detection algorithms. These datasets come from various domains and they differ in size, outlier percentage, and dimensionality, offering a broad evaluation spectrum.

Dataset	Size	Dimensions	Outliers	Percentage
b-cancer - <a href="#">Lichman et al. (2013)</a>	367	30	10	2.72%
pen-global - <a href="#">Lichman et al. (2013)</a>	809	16	90	11.12%
letter - <a href="#">Lichman et al. (2013)</a>	1,600	32	100	6.25%
speech - <a href="#">Brümmer et al. (2012)</a>	3,686	400	61	1.65%
sattelite - <a href="#">Lichman et al. (2013)</a>	5,100	36	75	1.47%
per-local - <a href="#">Lichman et al. (2013)</a>	6,724	16	10	0.15%
ann-thyroid - <a href="#">Lichman et al. (2013)</a>	6,916	21	250	3.61%
shuttle - <a href="#">Lichman et al. (2013)</a>	46,464	9	878	1.89%
aloi - <a href="#">Geusebroek et al. (2005)</a>	50,000	27	1,508	3.02%
kdd99 - <a href="#">Lichman et al. (2013)</a>	620,098	29	1,052	0.17%

Table 4: Datasets for Benchmarking

## Appendix B. Evaluation of Computational Performance

We now review the computational performance of the algorithms that are most prevalent in unsupervised clustering and are featured in this work, both for global clustering, and outlier detection. GU-Clustering operates in two modes, a Global Clustering mode, and a Global Anomaly Detection mode. It is designed to be simple to implement, yet highly effective and efficient. Since the data are stored in contiguous blocks of memory, it allows to perform component-wise operations. Component-wise operations are computationally efficient, and allow for vectorized operations, rather than looping over each element of the data structure individually.

Algorithm	Run time complexity
K-means	$O(nKdi)$
GMM	$O(nKd^3)$
FMCD (RC)	$O(n\log(n))$
DBSCAN	$O(n\log(n))$
LOF	$O(n^2) + O(n\log(n))$
oSVM	$O(n^2)$
iForest	$O(n)$
GU-Clustering	$O(d)$

Table 5: Run Time Complexity of featured unsupervised clustering algorithms

Although there are various implementations of the algorithms that are presented above, we review the variants used for benchmarking our proposal. Generally, algorithms that use

covariance matrices to uncover the relationships between points tend to be slow, since a run time complexity of  $O(n^3)$  is required for computing the covariance matrix, and cannot be applied efficiently to large datasets. Despite that, it is worth mentioning that in the case of oSVM, the run time complexity reduces to  $O(n^2)$  when the contamination parameter is provided. FMCD (Robust Covariance) is more efficient since it uses a sample of the dataset for estimating the covariance matrix. Traditional K-means needs to sum the distance for  $d$  dimensions, for  $n$  points, and check for  $K$  clusters if the point can be included. Usually, it takes a number of  $i$  iterations to converge to the optimal result. Depending on the implementation of the algorithm this can be reduced to a time complexity of  $O(nKi)$ . iForest is the most efficient among the competing algorithms operating in linear time for the evaluation of each point. DBSCAN and LOF also operate in a polynomial run time complexity, for visiting every neighboring region and testing each point as a candidate for that region. LOF however, does default to a higher run time complexity, since it needs to compute the score for each data point and work out the upper/lower bounds for its minpts parameter. GU-Clustering is by far the most efficient algorithm since all its operations are performed component-wise, and requires only  $d$  iterations for fitting the model. In the evaluation process, the run time complexity is reduced to  $O(K)$  for the global clustering mode and  $O(1)$  for the outlier detection mode.

<b>Dataset</b>	<b>Robust Cov.</b>	<b>oSVM</b>	<b>iForest</b>	<b>LOF</b>	<b>GU-Clustering</b>
b-cancer	0.13s	0.00s	0.31s	0.02s	0.00s
pen-global	1.12s	0.02s	0.34s	0.02s	0.00s
letter	2.37s	0.06s	0.39s	0.12s	0.00s
speech	65.9s	0.46s	2.58s	0.54s	0.01s
sattelite	5.58s	0.13s	0.77s	0.63s	0.00s
pen-local	4.67s	0.09s	0.81s	1.24s	0.00s
annthyroid	1.18s	0.41s	0.80s	1.16s	0.00s
shuttle	21.9s	5.90s	3.61s	5.09s	0.01s
aloi	11.5s	22.4s	5.25s	36.7s	0.01s
kdd99	161s	0.79hrs	50.9s	2.13hrs	0.14s

Table 6: Computational performance comparison of the algorithms evaluated in section 4.2. The experiments were conducted on a Google Colab environment with an Intel Xeon(R) CPU (1 core @ 2.2 GHz) and 12 GB of memory. GU-Clustering is more frugal on resources and significantly more efficient in terms of computational time.

## Appendix C. Visual Inspection with Prominent Anomaly Detection Algorithms

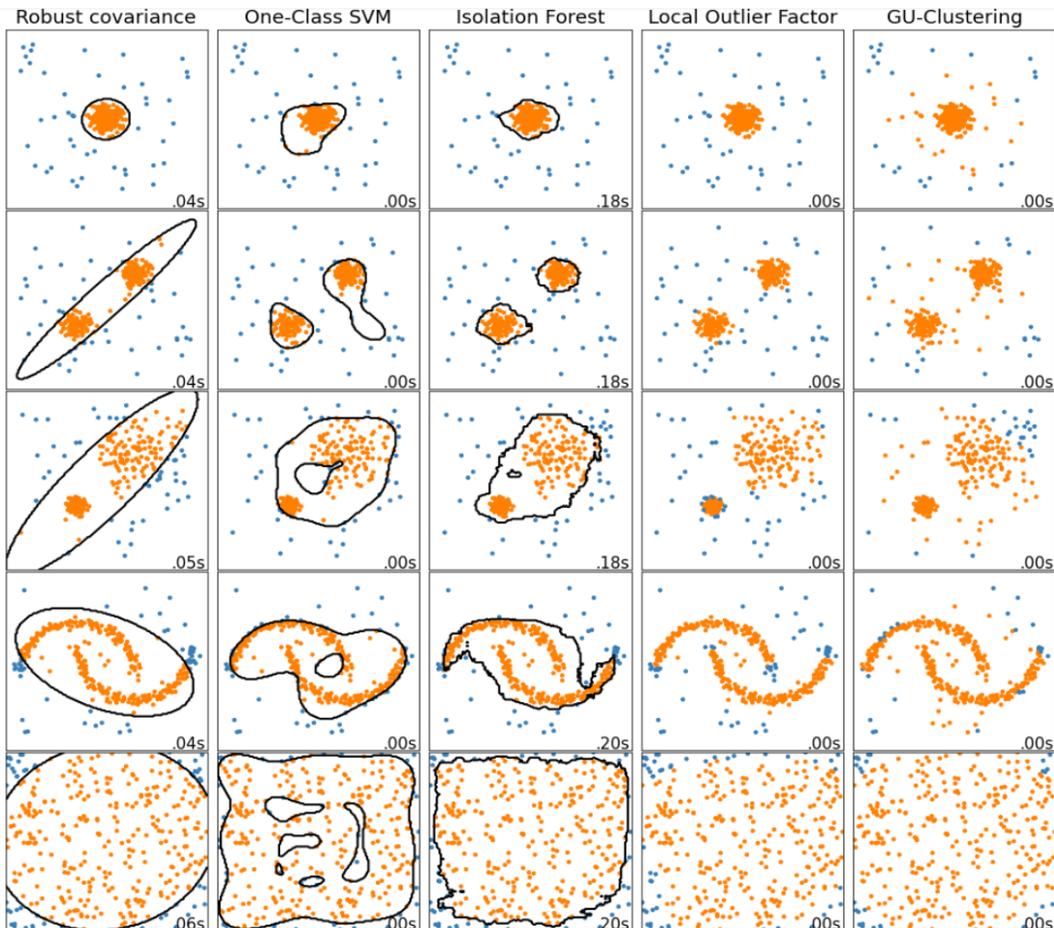


Figure 2: This application is an extension of scikit-learn’s section of comparing outliers for anomaly detection [Pedregosa et al. \(2011\)](#). The datasets contain one or two modes (regions of high density) to illustrate the ability of the algorithms to cope with multimodal data. For each dataset, 15% of samples are generated as random uniform noise. This proportion is the value given to the  $\nu$  parameter of the oSVM and the contamination parameter of the other outlier detection algorithms. The decision boundaries between inliers and outliers are displayed in black except for LOF and GU-Clustering as they do not have a predict method to be applied to new data. All algorithms parameters were hand-picked by the authors for achieving the best results. GU-Clustering effectively detects global outliers, as seen in the figure above, providing a new perspective to outlier detection. Similar to the results from the previous sections, the outliers detected focus on the global position of the points rather than outliers belonging to specific densities. Similar to LOF, GU-Clustering does not need to be fitted to a dataset before classifying new points, thus offering the possibility of detecting outliers in real time. Furthermore, GU-Clustering does not require prior knowledge of the contamination fraction and can achieve state-of-the-art performance using a default value for  $q$ .