

# Global density clustering

*Scott Cohen*

This paper presents Global Density Clustering, (GDC), an algorithm that has several major advantages over the most popular existing clustering algorithms: (1) No parameters are chosen at the outset of the function; rather, the user can control the desired resolution as clustering proceeds. (2) GDC is efficient enough to work on a large dataset even when there are a sizable number of features. It is  $O(MN \log N)$  where  $M$  is the number of features, i.e. the dimension, and  $N$  is the number of data points, i.e. the dataset size. It is suitable for big data. (3) GDC has the advantage of the powerful and intuitive definition of clusters as: points within a cluster are closer than distance  $dist$  to their nearest neighbor in the cluster ( $dist$  is not picked at the outset but rather that is chosen as the algorithm is progressing) and all points outside the cluster are further than  $dist$  from any point in the cluster. (4) GDC supports variable density without the plethora of special data structures such as HDBSCAN needs. (5) Other advantages are described. An essential reason that GDC has these advantages is it searches for and considers points whose nearest neighbor are furthest apart before searching for those that are closer together. It is a top-down or “global” consideration of distances that other density algorithms do from a bottom-up or “local” view. Other novel approaches to the main problems of clustering such as noisy backgrounds are described.

## I. INTRODUCTION

All clustering algorithms have issues. According to the Impossibility Theorem for clustering, there is no best clustering algorithm [1]. Multiple valid meaningful clusters can be constructed from the same data. See section 3.2 in [2]. This paper, however, proposes an algorithm that has many advantages compared to the most popular existing choices.

[3] is a comprehensive survey of clustering algorithms with a table that tabulates the strengths of many clustering algorithm on 7 aspects: time complexity; scalability; handles large-scale data; handles high dimensional data; suitable shape of cluster; sensitive to sequence of inputting data; sensitive to noise and outliers. Large-scale data includes the issues of the 4 Vs: volume, velocity, variety, and incomplete veracity.

There are also, however, other issues that arise for some but not all clustering algorithms. Some of these are: whether parameters need to be chosen before the algorithm runs especially if the parameters are difficult to choose and the results are sensitive to that choice; whether the algorithm is easily explainable especially in regulated fields since a person oversees the result; how well does the algorithm handle streaming of a large update of data; does it handle categorical data; does it have unusual data structures that take time to construct, interpret, and explain.

The following comparisons touch on the above issues and are organized by discussing them in the context of the most well-known and popular current clustering algorithms.

There are many disadvantages of k-means clustering that do not occur with GDC. With GDC:

- There is no equivalent of  $k$ , let alone an arbitrary choice of  $k$ . In fact, no parameters are chosen at the outset of the algorithm; as the algorithm progresses, the user has the option to stop it when user likes the clusters found thus far. The user decides when the resolution is right. A

- suggestion to stop is made when nearest neighbors get much closer together than before. This is the GDC solution to one of the most difficult problems in clustering. See section 3.3 in [2].
- Many kinds of clusters are captured that k-means misses even in the simplest of situations such as in 2-dim let alone in higher dimensions. See section 3.1, figure 5 in [2] for one example.
  - There is no need to worry whether the clusters are roughly spherical as with k-means. See table 22 in [3].
  - All this is efficiently supported in an arbitrary number of dimensions.
  - It is suitable for big data with or without any sampling. The entire (big) database can be efficiently clustered.
  - There is no concern about sensitivity to an initial choice of centroids.

Perhaps the biggest advantage of k-means is that it is fast:  $O(tMN)$  where  $M$  is the number of features,  $N$  is the dataset size, and  $t$  is the number of iterations of the k-means algorithm. Algorithmically, global density clustering is not far behind if  $t$  is at least slightly bigger than  $\log N$  which is not unreasonable. Global density clustering is  $O(MN \log N)$  so the comparison depends on  $t$  versus  $\log N$ . Sort routines even on big data are so highly optimized and  $\log N$  even for big data is probably somewhat smaller than 40. So, the comparison of runtimes is reasonable.

Although global density clustering is, as its name suggests, a density algorithm, it does not suffer from some of the disadvantages of other density algorithms such as DBSCAN, OPTICS, and HDBSCAN. With GDC:

- There is no need or use for parameters min pts or epsilon let alone any arbitrary choice of them as with DBSCAN and OPTICS.
- Clusters are caught that would be missed because of the arbitrary choice of epsilon as with DBSCAN and OPTICS, especially in higher dimensions.
- The worst-case runtime is better than DBSCAN, OPTICS, HDBSCAN which are worse than  $O(N \log N)$ . See [4]. It is  $O(MN \log N)$  where  $M$  is the number of features (i.e. dimensions) and  $N$  is the number of points in  $M$ -dim (i.e. the dataset size). This is the difference between practical and for the worst-case situation prohibitively intolerable as with DBSCAN, OPTICS and HDBSCAN.
- There is no need or use for a dendrogram as with OPTICS.
- There is no need or use for a minimum spanning tree, cluster hierarchy, condensed tree or any other special purpose structures that take time to construct, interpret, and explain as with HDBSCAN.
- We have easy explanations of what the algorithm is doing unlike HDBSCAN.

Density algorithms all have a major benefit in their intuitive definition of the clusters that are uncovered, and the arbitrary cluster shapes supported. Likewise, GDC has those benefits. GDC definition of a cluster is that points are within a cluster if they are closer than distance  $dist$  to their nearest neighbor in the cluster and that all points outside the cluster are further than  $dist$  from any point in the cluster. (The distance  $dist$  is not picked at the outset but rather that is chosen as the algorithm is progressing.) In addition, GDC has the advantages described above; they are significant.

Hierarchical clustering has some serious performance disadvantages. Regarding GDC:

- GDC is exponentially faster. It is  $O(MN \log N)$  where  $M$  is the number of features (i.e. dimensions) and  $N$  is the number of points in  $M$ -dim (i.e. the dataset size).
- GDC only needs two calls to sort the data per feature and one of them nearly always exists anyway because it'll be a primary or secondary index.

Until GDC, the benefits of hierarchical clustering were more theoretical than practical because of its runtime performance.

In general, choosing such parameters such as  $k$ , epsilon, min pts, etc. is largely a guessing game. At best they indicate some a priori knowledge such as there are 10 digits for the MNIST dataset, so we say for  $k$ -means that there are 10 clusters at the outset of the algorithm. But choosing these parameters affect the resulting clustering which defeats the exploratory advantage of clustering algorithms. We are affecting the result before we know what the best choice for clusters is.

In general, it would be useful to avoid all these a priori choices which can affect the outcome before we start. It is valuable to be easy to explain. It is also very valuable to be usable for big data.

## II. GLOBAL DENSITY CLUSTERING

The approach taken in this paper is, first, completely solve the problem when each data point is a scalar i.e. just a single real number and then, second, build up the solution to when each data point is a  $M$  dimensional point where  $M$  is the number of features of the data.

The complete solution of the 1-dimensional data point used here is based on the observation that the clusters along a line with the  $N$  points are simply bounded by the largest of the  $N-1$  intervals that are in descending order. This view occurs as a result of looking at the global picture and drilling top down. Although this is a density algorithm in that we look for pockets of points close to each other (usually a local view) we use top down processing (usually a global view) to find the dense pocket of data points. Hence, the name of this algorithm.

Let's say that when we pick the largest interval in one dimension, it is `largest_interval_remaining`. Then all, say,  $k < N$  data points to the left of `largest_interval_remaining` are in one group of clusters. Likewise, all  $N-k$  data points to the right of `largest_interval_remaining` are in a different group of clusters. Every point in the left group of clusters is further than  $|\text{largest\_interval\_remaining}|$  from every point in the right group of clusters. This procedure then may be, but not necessarily, repeated for one or both the left group and the right group for `second_largest_interval_remaining`, and so on.

Now we discuss building up the solution from 1-dim to where each data point is an  $M$  dimensional point. The approach relies on a few observations:

- First, if two points are in the same cluster then they should be in the same cluster in each projection onto an axis (even though the reverse is not necessarily true);
- Second, in  $M$ -dim space, if the distance is  $\text{dist}$  between two points then the projection onto at least one axis is  $\geq \text{dist}/M$ ; otherwise the  $M$  lengths would add up to less than  $\text{dist}$  when in L1 norm which is a contradiction.
- Third, the minimum distance between points of different clusters is, in L1 norm,  $\sum |\text{largest\_interval\_remaining}|$  whereas the maximum distance between two nearest neighbors is  $\sum |\text{second\_largest\_interval\_remaining}|$  which can be much smaller.

Distance is in L1 norm for convenience of analysis; similar situations occur for L2 norm as well as other norms. Also, the axes could be any  $M$  independent vectors although we simply choose the orthogonal coordinate axes so that their projection onto that axis is just selecting its coordinate of the  $M$  coordinates.

An essential reason that GDC has these advantages is it searches for and considers points whose nearest neighbor is furthest apart before searching for those that are closer together. It is a top-down or "global" consideration of distances that other density algorithms do from a bottom-up or "local" view.

A careful reading of the above description indicates that GDC is a divisive hierarchical algorithm as well as a density algorithm. As a result of the top-down consideration of clusters, variable density of the clusters is handled easily: clusters defined earlier in the top-down processing have lower density than those clusters that are defined later.

GDC handles streaming data updates by collecting as much of the new data being streamed into memory of a single computer where it can sort it fast. Then GDC merges the sorted updates into the existing data which is also sorted. This is simply another merge iteration of the mergesort used by GDC. During this merge, nearest neighbor distances are updated so that the top-down processing of defining clusters can be repeated. But since the sorting is the reason that time complexity is  $O(MN \log N)$  and it is done as just mentioned, it will be practical for GDC to repeat defining clusters based on the new information of nearest neighbors. In this way, streaming data fits naturally into the GDC algorithm.

Noisy data is particularly problematic for many algorithms including GDC. Distinguishing outliers from noise is very difficult. Another example of a problem is where noise causes two clusters that should be separate to be confused with a single cluster. DBMAC and DBMAC II [5] approaches the problem of a noisy background (even 80%) with multi-scale analysis. GDC considers two alternatives to multi-scale analysis.

Both approaches by GDC to handle noise take advantage of the step where GDC projects the dataset onto each 1-dimensional axis. At that time, GDC can insert code to deal with noise in 1-dim space. The heavy-handed code to deal with noise removal in 1-dim space is Fast Fourier Transforms (FFT). In particular,

1. FFT the 1-dim projections of the dataset,
2. Remove higher frequencies just as if denoising an audio signal
3. Do the inverse FFT back from frequency space to the 1-dim projection
4. Remove the data points in the 1-dim projection that are not in the result of the inverse FFT
5. Remove the M-dimensional points that projected onto those points in step 4.
6. Continue with GDC processing.

FFT is  $O(N \log N)$  time complexity so the impact on runtime depends on how many axes, one or more, the above steps are done.

A less heavy-handed approach with much faster time complexity is to simply subtract the average density from the 1-dim projections of the dataset, In particular,

1. Set `avg_density` to the range of values on the 1-dim divided by N
2. Pick at random, uniform distribution, points on this axis with a probability of `avg_density`.
3. Remove the points in 1-dim
4. Remove the M-dimensional points that projected onto those points in step 3.
5. Continue with GDC processing

For this less heavy-handed approach, it is important to note that no assumption of constant density within the cluster is assumed. Cluster density is irrelevant. Rather the assumption is of homogenous density of the average amount of background noise. This might be a more realistic restriction. Even that assumption might be ameliorated by treating different regions of the dataset separately.

The only step that is not  $O(MN)$  in all the above for GDC is a sort of the dataset on each feature and possibly a FFT on one or more 1-dim axes; even for the huge datasets of big data this is not prohibitive. Therefore, it is not an order  $O(N^2)$  algorithm. In the code for this paper, mergesort is used because of its excellent performance on big data. All other processing of data is done with sequential passes over the

data so only sequential disk IO is necessary and that results in the best performance for runtime especially for big data.

### III. WHAT IS A COMPLETE CLUSTERING SOLUTION?

We define a complete clustering solution in any dimensional space to be a partition of data points into clusters of 1 or more data points that satisfy two properties:

For some number  $dist$ , where  $v1, v2$  are two data points (i.e. vectors), for all  $v1$  in one cluster we can find a partition into one or more sets of data points where

(\*) if there is more than one data point  $v1, vk$  in the cluster then there exists a sequence of points  $v_i$  starting at  $v1$  and ending at  $vk$  all in that cluster,  $|v_{i+1} - v_i| < dist$ ;

(\*\*) all  $v2$  in a different cluster,  $|v1-v2| \geq dist$

Note that this paper uses the L1 norm because it is easier to analyze in this situation than the L2 norm. So if  $dist$  is the distance between  $v1$  and  $v2$ ,  $dist = \sum |v1[axis\_ndx] - v2[axis\_ndx]|$

There is a choice of clusters that satisfies properties (\*) and (\*\*) and for a given  $dist$  it is unique. The argument is for each point A: add all points to the cluster that are within distance  $dist$  of A; keep repeating the previous step for each point in the cluster until either all points in the entire dataset are in the cluster or not. This satisfies property (\*). If not all points are in the cluster containing point A then all the remaining points are distance  $\geq dist$  from all points in the cluster which is property (\*\*). Now we can repeat the above steps by finding a cluster that is distinct from the cluster containing point A.

In order to get a little more understanding, consider that as  $dist$  increases, clusters get larger and there are fewer clusters. The largest  $dist$  can be is the diameter of the dataset in  $R^M$ . As  $dist$  decreases the clusters get smaller and eventually there are N clusters where each is an isolated point like an outlier. The smallest  $dist$  can be is the smallest interval which can be close to zero.

So, letting  $dist$  start out as the diameter or larger of the dataset in  $R^M$ , we start top-down from a global view with one cluster and as  $dist$  gets smaller when we set it to the largest remaining distance between two points. We iteratively apply the first step to either the left group of clusters or to the right group of clusters.

We continue decreasing  $dist$  until the next interval is a dramatic drop in size (the ratio is tiny) or until the user is satisfied with the clusters found thus far. A dramatic drop in size might indicate that any more clusters will already be dense enough to form a single cluster. But the user can override this and continue clustering.

One could think of (\*) as a local statement about nearest neighbors in just one cluster and just about its insides. On the other hand, (\*\*) is basically a global statement about points that are not nearby but far apart in two clusters i.e. global objects.

Picking the largest interval on an axis may in practice initially pick out the outliers and then will probably start to divide the data points roughly in half. If so then there will be roughly the number of outliers plus  $\log(N)$  repetitions which helps make this an even faster algorithm. Regardless, the worst case runtime is only  $O(MN \log N)$ .

### IV. PSEUDO-CODE FOR COMPLETE SOLUTION IN 1-DIMENSION

First, we show the python pseudo-code for global density clustering in 1-dimension. It is a complete solution.

Start with a  $\text{dist} > \text{diameter}$ , a cluster containing the entire dataset, and a single axis line.

$\text{sorted\_N\_points} = \text{sort the } N \text{ 1-dim data points that are on that axis}$

$\text{largest\_interval\_array} = \text{sort-in-descending-order}(\text{diff-2-consecutive}(\text{sorted\_N\_points}))$

for  $\text{interval\_ndx}$  in  $\text{range}(N)$ :

$\text{largest\_interval\_remaining} = \text{largest\_interval\_array}[\text{interval\_ndx}]$

    Say  $\text{largest\_interval\_remaining}$  is bounded by data points  $I_L$  on the left and  $I_R$  on the right.

    If  $I_L$  and  $I_R$  both belong to the same cluster

        and that cluster is bounded by data points  $C_L$  on the left and  $C_R$  on the right:

    then replace that cluster with the two clusters;

        one, data points in  $[C_L, I_L]$  and, two, data points in  $[I_R, C_R]$

    If  $\text{largest\_interval\_array}[\text{interval\_ndx}+1]$  is much smaller than  $\text{largest\_interval\_remaining}$ :

    then ask the user whether to quit

    reduce  $\text{dist}$  to  $\text{dist} = \text{largest\_interval\_array}[\text{interval\_ndx}]$

/\* when done the clusters satisfy (\*) and (\*\*) \*/

## V. THE APPROACH FOR MORE THAN 1-DIM DATA POINTS

For simplicity, let's start with  $M=2$ . Let 1-dim point  $A_x$  be the 2-dim point  $A$  projected onto the x-axis and 1-dim  $A_y$  be the point  $A$  projected onto the y-axis. Likewise, for 2-dim point  $B$  and 1-dim points  $B_x$  and  $B_y$ . Then either (remember that  $M=2$  for now)

$$|B_x - A_x| > |B - A|/M$$

where the left side of the inequality is distance along a 1-dimensional interval bounded by 1-dim points  $A_x$  and  $B_x$  and the right side of the inequality is an  $M$ -dim distance along a  $M$ -dimensional line segment bounded by  $M$ -dim points  $A$  and  $B$ ; or

$$|B_y - A_y| > |B - A|/M$$

for the 1-dim points  $A_y$  and  $B_y$ . This either/or inequalities is because of the L1 norm where

$$|B - A| = |B_x - A_x| + |B_y - A_y|$$

Define

$K^{\text{dist}}$  is the cluster in  $M$ -dim containing  $A$  for distance  $\text{dist}$ ,

$K_x^{\text{dist}}$  is the cluster of the projection of  $K^{\text{dist}}$  onto the x-axis, in 1-dim containing  $A_x$  for distance  $\text{dist}$ .

By the definition of  $K_x^{\text{dist}}$ , for every axis  $x$ ,

(\*\*) if  $A_x \in K_x^{\text{dist}}$  and  $B_x$  is not then  $|A_x - B_x| > \text{dist}$ .

(\*) if  $Ax$  and  $Bx \in K_x^{\text{dist}}$  and  $Ax$  and  $Bx$  are nearest neighbors then  $|Ax - Bx| < r_x * \text{dist}$  where  $r_x = (\text{second\_largest\_interval\_remaining} / \text{largest\_interval\_remaining})$ .

Notice that, for each axis,  $r_x \leq 1$  because  $\text{second\_largest}$  is always smaller or equal to  $\text{largest}$ .




Figure 1.

Just in case the figure is not clear:

$$Ax + \text{dist} = Bx = B'x; B''x = Bx + 2 * \text{dist}; Ay = By = B''y; B'y = By + 2 * \text{dist}$$

The entire dataset in figure 1 is four M-dim points where  $M=2$ ,  $\{A, B, B', B''\}$ .  $K_x^{\text{dist}}$  is that 1-dim cluster that contains  $Ax$ ;  $K_y^{\text{dist}}$  is that 1-dim cluster that contains  $Ay$ ; 1-dim clusters on x-axis are  $\{Ax, Bx\}$ ,  $\{B''\}$  and on y-axis are  $\{Ay\}$ ,  $\{B'y\}$ ; the M-dim clusters are  $\{A, B\}$ ,  $\{B'\}$ ,  $\{B''\}$ ;  $\text{inverse}(K_x^{\text{dist}}) = \{A, B, B'\}$ ;  $\text{inverse}(K_y^{\text{dist}}) = \{A, B, B''\}$

For simplicity say  $M = 2$ . Consider the set of dataset points  $K'$  in M-dim space

$$K' = \{\cap \text{inverse}(K_x^{\text{dist}})\}$$

where  $\text{inverse}()$  takes a set of 1-dim points  $K_x^{\text{dist}}$  and maps them to the set of M-dim points that project onto  $K_x^{\text{dist}}$ .

The subset  $K'$  is in fact a solution: it satisfies (\*\*) because

If  $A \in \{\cap \text{inverse}(K_x^{\text{dist}})\}$  where the intersection is over every axis and  $B$  is not in  $\{\cap \text{inverse}(K_x^{\text{dist}})\}$  then

$$B \in \{\cup \{\text{inverse}(K_x^{\text{dist}})^{\text{complement}}\}$$

for at least one axis  $x$ ,  $B \in \{\text{inverse}(K_x^{\text{dist}})^{\text{complement}}\}$ .

$$B \in \{\text{set of all M-dim points } Z \text{ s.t. } |Zx - Ax| < \text{dist}\}^{\text{complement}}.$$

For at least one axis  $x$ ,  $|Bx - Ax| > \text{dist}$

which implies  $|B - A| > \text{dist}$ . Also,  $K'$  satisfies (\*) because

for every axis  $x$ ,  $|Bx - Ax| < \text{dist}$  where  $\text{dist}$  will be some  $\text{largest\_interval\_remaining}$ ; so,

$|Bx - Ax| < r_x * \text{dist}$  where  $r_x * \text{dist}$  will be the  $\text{second\_largest\_interval\_remaining}$  and summing

$$\sum |Bx - Ax| < \sum r_x * \text{dist}$$

$$|B - A| < \text{dist}$$

if

$$\sum r_x < 1.$$

But this is a reasonable inequality to expect at some time as we reduce distance `dist` in descending order of `largest_interval_remaining`. In English, it says that if the density decreases enough on all the axes (even if, most likely, by different amounts) then we can make a cluster for that distance that is a complete solution.

If distance `dist` does not satisfy  $\sum r_x < 1$  then an alternative approach is to increase `dist` slightly to `dist'` so that  $K_x^{\text{dist}'}$  does satisfy (\*).

Note that when we include a new  $M$ -dim point the maximum distance between nearest neighbors in the same cluster might decrease along one or more axes and thus in  $M$ -dim space as well. Yet the partition cluster that contains point  $A$  will still obey (\*). Similarly, the minimum distance between two clusters may increase along one or more axes and thus increase in  $M$ -dim space. Yet the partition into clusters will still obey (\*\*).

An outline of the algorithm is in the section below. This approach is a complete solution that always satisfies (\*) and (\*\*). The algorithm is  $O(MN \log N)$ .

## VI. PSEUDO-CODE FOR MORE THAN 1-DIMENSION

Now, we show the python pseudo-code for global density clustering in  $M$ -dimensions and  $M$  axes `axes[0], ... axes[M-1]`. Note how it builds on the solution in 1-dimension axis for each axis.

```
for axis_ndx in range(M):
```

```
    Get all clusters on the line axes [axis_ndx] for distance dist that is one of the largest intervals using the 1-dimension solution.
```

```
    Say the clusters are D1, D2, ... Dn' where n' is
```

```
    largest_interval_remaining[axis_ndx][n' - 1] = dist >
largest_interval_remaining[axis_ndx][n']
```

```
    sum_all_second_largest += largest_interval_remaining[axis_ndx][n']
    which happens to be the second_largest_interval_remaining[axis_ndx]
```

```
    Replace each cluster Ci with the group of clusters {{Ci ∩ D1}, ..., {Ci ∩ Dn'}}
```

```
    reduce dist to dist= largest_interval_array[interval_ndx]
```

```
    /* when done the clusters satisfy (*) and (**) */
```

```
if sum_all_second_largest < dist:
```

```
    adjust dist and repeat the above loop until either this inequality holds or
    we decide there is no clear-cut further division into clusters.
```

When implemented efficiently, this algorithm is  $O(MN \log N)$ .



## VII. RESULTS AND COMPARISON WITH OTHER CLUSTERING ALGORITHMS

We investigate how well global density clustering does with MNIST images of digits: (1) suggesting the correct number of clusters, ten for the digits; (2) accurately predicting the label of a test digit by seeing which cluster we assign the test digit to and seeing the label that we assigned to that cluster during the training phase. The label assigned to that cluster was the plurality of the labels of all the digits in that cluster during training phase. We then do the same with k-means and DBSCAN and compare the results. We can even compare the results to the accuracy of supervised learning models such as kNN, XGBOOST, etc. All this is work that is in progress.

## VIII. FUTURE INVESTIGATIONS

An important action is to try different datasets. MNIST is the “hello world” dataset for machine learning algorithms but clustering has more practical uses such as in market segmentation. Such datasets would provide real-world use and comparison of the global density algorithm.

The second item to do is to compare against other “all nearest neighbor” algorithms in addition to the above comparison with k-means and DBSCAN. For example, [6] Amitabh Varshney, is orders of magnitude faster than alternatives. It is important to see how well global density does in comparison especially on big data.

Of course, there are many other directions to explore. For example, exploring norms other than L1 would be nice. More important is exploring Hamming distance for categorical clustering. An important study will be the impact of GDC ideas on other cluster algorithms and cluster methods (a distinction from [5]). It will be informative to see how it fits in with the clustering of cluster algorithms by their clustering of 12 datasets. See section 3.5 in [2].

But most important is that future investigation must try clustering big data to analyze its performance. We would learn the limits of M and N for not being prohibitively costly in runtime performance. Large scale clustering is where we cluster millions of data points with thousands of features. Gene clustering, clustering of earth science data and other extremely valuable applications have an M of order 100 and an N of order 100,000 and requires support for large scale clustering. See section 4.3 in [2]. This is where GDC may have a particularly significant impact.

## IX. REFERENCES

- [1] Kleinberg J An impossibility theorem for clustering NIPS 15 MIT Press 463-470 (2002).
- [2] Jain, A, Data clustering: 50 years beyond K-means, *Pattern Recognition Letters* Volume 31, Issue 8, 1 June 2010, Pages 651-666 <https://doi.org/10.1016/j.patrec.2009.09.011>
- [3] Xu, D., Tian, Y. A Comprehensive Survey of Clustering Algorithms. *Ann. Data. Sci.* **2**, 165–193 (2015). <https://doi.org/10.1007/s40745-015-0040-1>
- [4] Schubert E., Sander J., Ester M., Kriegel H. P. DBSCAN Revisited, Revisited: Why and How You Should (Still) Use DBSCAN ACM Transactions on Database Systems, Vol. 42, No. 3, Article 19. 2017.
- [5] T. Zhang and B. Yuan, "Density-Based Multiscale Analysis for Clustering in Strong Noise Settings With Varying Densities," in *IEEE Access*, vol. 6, pp. 25861-25873, 2018.

- [6] Sankaranarayanan, J., Samet, H. A fast all nearest neighbor algorithm for applications involving large point-clouds *Computers & Graphics*, **2**, 157-174 (2007).
- [7] Xu R, Wunsch D Survey of clustering algorithms. *IEEE Trans Neural Netw* 16:645–678 (2005).