



**HAL**  
open science

## A distributed approximate nearest neighbors algorithm for efficient large scale mean shift clustering

Gaël Beck, Tarn Duong, Mustapha Lebbah, Hanane Azzag, Christophe Cérin

### ► To cite this version:

Gaël Beck, Tarn Duong, Mustapha Lebbah, Hanane Azzag, Christophe Cérin. A distributed approximate nearest neighbors algorithm for efficient large scale mean shift clustering. *Journal of Parallel and Distributed Computing*, 2019, 134, pp.128-139. 10.1016/j.jpdc.2019.07.015 . hal-04456309

**HAL Id: hal-04456309**

**<https://hal.science/hal-04456309>**

Submitted on 15 Mar 2024

**HAL** is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

# A Distributed and Approximated Nearest Neighbors Algorithm for an Efficient Large Scale Mean Shift Clustering

Gaël Beck<sup>1</sup>, Tarn Duong<sup>1</sup>, Mustapha Lebbah<sup>1</sup>, Hanane Azzag<sup>1</sup>, and  
Christophe Cérin<sup>1</sup>

<sup>1</sup>Computer Science Laboratory of Paris North (LIPN, CNRS UMR 7030), University of Paris 13, F-93430 Villetaneuse France

## Abstract

In this paper we target the class of modal clustering methods where clusters are defined in terms of the local modes of the probability density function which generates the data. The most well-known modal clustering method is the  $k$ -means clustering. Mean Shift clustering is a generalization of the  $k$ -means clustering which computes arbitrarily shaped clusters as defined as the basins of attraction to the local modes created by the density gradient ascent paths. Despite its potential, the Mean Shift approach is a computationally expensive method for unsupervised learning. Thus, we introduce two contributions aiming to provide clustering algorithms with a linear time complexity, as opposed to the quadratic time complexity for the exact Mean Shift clustering. Firstly we propose a scalable procedure to approximate the density gradient ascent. Second, our proposed scalable cluster labeling technique is presented. Both propositions are based on Locality Sensitive Hashing (LSH) to approximate nearest neighbors. These two techniques may be used for moderate sized datasets. Furthermore, we show that using our proposed approximations of the density gradient ascent as a pre-processing step in other clustering methods can also improve dedicated classification metrics. For the latter, a distributed implementation, written for the Spark/Scala ecosystem is proposed. For all these considered clustering methods, we present experimental results illustrating their labeling accuracy and their potential to solve concrete problems.

**Index terms**— Clustering, Gradient Ascent, Nearest neighbors, Spark

## 1 Introduction

The goal of clustering or unsupervised learning is to assign cluster membership to unlabeled candidate points where the number and location of these clusters

are unknown. Current Mean Shift clustering [1, 2, 3] algorithms contains computational bottlenecks with both kernel and nearest neighbor approaches: the former is due to the exact evaluation of the kernel function, and the latter due to the exact nearest neighbor searches. We propose a new algorithm which resolves the computational inefficiencies of the nearest neighbor Mean Shift by using Locality Sensitive Hashing (LSH) [4, 5, 6] for approximate nearest neighbor searches to replace the exact nearest neighbor calculations in the density gradient ascent and in the cluster labeling stages. Compared to kernel approaches to Mean Shift clustering, which are  $O(n^2)$  where  $n$  is the size of the dataset, our nearest neighbors approach enables a scalable implementation of the gradient ascent and cluster labeling which are both  $O(n)$ .

Furthermore, existing programming paradigms for dealing with parallelism, such as MapReduce [7] and Message Passing Interface (MPI) [8], have been demonstrated the best practical choices for implementing these clustering algorithms. MapReduce paradigm becomes popular and suited for data already stored on a distributed file system, which offers data replication, as well as the ability to execute computations, locally on each data node. Thus, we implement this approximate nearest neighbour Mean shift clustering algorithm on a distributed Apache Spark/Scala framework [9], which allows us to carry out clustering on Big datasets.

The organization of the paper is as follows. We review related works in Section 2. In Section 3 we briefly introduce the problem prior to elaborating our algorithmic contributions for the gradient ascent and cluster labeling. Section 5 is related to the experiments conducted on the Grid'5000 testbed where we examine the role of the key tuning parameters for both accelerating the respond time and measuring the quality of the answer Concluding remarks follow in Section 6.

## 2 Related works

For comprehensive reviews of clustering, see for example the monographs as in [10, 11]. For our purposes, we focus on the principal scalable clustering algorithm, i.e. on the well-known  $k$ -means algorithm [12]. This algorithm has the advantage of having a single parameter  $k$  which stands for the number of desired clusters. The algorithm works by moving  $k$  prototypes towards the centroid formed by their closest data points. This process is repeated iteratively until the intra-class variance (the sum of squared distances from each data point within a cluster to its corresponding prototype) is minimized. DBScan [13] is a well-known density based algorithm. It takes two parameters,  $\epsilon$  which defines the radius of the hypersphere and  $minPts$  which is the minimum number of points above which a corresponding hypersphere is considered to be sufficiently dense. Each time the density threshold is

reached, the data points in the same hypersphere belong to the same cluster, and the process is extended to include more data points until the data density falls under the threshold determined by  $\varepsilon$  and *minPts*. The remaining data points are considered to be noise. One notable advantage of this algorithm is its ability to automatically detect the number of clusters with arbitrary shape. But it remains difficult to tune it correctly. Like DBScan, Mean Shift is a density based algorithm and can detect automatically the number of clusters with arbitrary shape. Most studies on the Mean Shift clustering have focused on the kernel versions, e.g. [14, 15, 16]. The latter authors compared Gaussian, Cauchy and generalized Epanechnikov kernels to study the behaviour of tuning parameters of Mean Shift clustering.

### 3 Mean Shift clustering

The Mean Shift algorithm and its variants consist in two major steps as described in Algorithm 1. The first important step (the density gradient ascent) is generally the most computationally intensive. This gradient ascent can be computed in different ways, such as with kernel functions or as we propose in this paper, nearest neighbors. The second step is the cluster labeling phase where we use the result from the first step to assign cluster labels to the original data points.

---

#### Algorithm 1 Mean Shift principle

---

**Input:** points  $\{\mathbf{x}_1, \dots, \mathbf{x}_m\}$ ,  
**Output:** cluster label  $\{\tilde{c}(\mathbf{x}_1), \dots, \tilde{c}(\mathbf{x}_m)\}$   
**Step 1:** Density gradient ascent;  
**Step 2:** Cluster labeling;

---

#### 3.1 Density gradient ascent

The Mean Shift method for a  $d$ -dimensional point  $\mathbf{x}$ , generates a sequence of points  $\{\mathbf{x}_0, \mathbf{x}_1, \dots\}$  which follows the gradient density ascent paths using the recurrence relation

$$\mathbf{x}_{j+1} = \frac{1}{k} \sum_{\mathbf{X}_i \in k\text{-nn}(\mathbf{x}_j)} \mathbf{X}_i \quad (1)$$

where  $\mathbf{X}_1, \dots, \mathbf{X}_n$  is a random sample drawn from a common density  $f$  and the  $k$  nearest neighbors of  $\mathbf{x}$  are  $k\text{-nn}(\mathbf{x}) = \{\mathbf{X}_i : \|\mathbf{x} - \mathbf{X}_i\| \leq \delta_{(k)}(\mathbf{x})\}$  with  $\delta_{(k)}(\mathbf{x})$  is the  $k$ -th nearest neighbor distance, and  $\mathbf{x}_0 = \mathbf{x}$ . Eq. (1) gives the Mean Shift method its name since the current iterate  $\mathbf{x}_j$  is shifted to the sample mean of its  $k$  nearest neighbors for the next iterate  $\mathbf{x}_{j+1}$ . The gradient ascent paths towards the local modes produced by Eq. (1) form

the basis of Algorithm 2, our nearest neighbor Mean Shift gradient ascent (NNGA).

The inputs to the NNGA are the data sample  $\mathbf{X}_1, \dots, \mathbf{X}_n$  and the candidate points  $\mathbf{x}_1, \dots, \mathbf{x}_m$ , which we want to cluster (these can be the same as  $\mathbf{X}_1, \dots, \mathbf{X}_n$ , but this is not required); and the tuning parameters: the number of nearest neighbors  $k$ , the tolerance under which subsequent iterations in the Mean Shift update are considered to be convergent  $\varepsilon_1$ , the maximum number of iterations  $j_{\max}$ .

---

**Algorithm 2** NNGA – Nearest Neighbor Gradient Ascent with exact  $k$ -nn

---

**Input:**  $\{\mathbf{X}_1, \dots, \mathbf{X}_n\}, \{\mathbf{x}_1, \dots, \mathbf{x}_m\}, k, \varepsilon_1, j_{\max}$   
**Output:**  $\{\mathbf{x}_1^*, \dots, \mathbf{x}_m^*\}$

- 1: Compute similarity matrix and sort each row;
- 2: **for**  $\ell := 1$  to  $m$  **do**
- 3:    $j := 0; \mathbf{x}_{\ell,0} := \mathbf{x}_\ell;$
- /\* Search for  $k$ -nn based on similarity matrix \*/
- 4:    $\mathbf{x}_{\ell,1} :=$  mean of  $k$ -nn( $\mathbf{x}_{\ell,0}$ );
- 5:   **while**  $\|\mathbf{x}_{\ell,j+1}, \mathbf{x}_{\ell,j}\| > \varepsilon_1$  **or**  $j < j_{\max}$  **do**
- 6:      $j := j + 1;$
- 7:      $\mathbf{x}_{\ell,j+1} :=$  mean of  $k$ -nn( $\mathbf{x}_{\ell,j}$ );
- 8:    $\mathbf{x}_\ell^* := \mathbf{x}_{\ell,j};$

---

The classical version of the NNGA introduced in Algorithm 2 requires, for each candidate point that we compute, the distance to all other data points, from which the mean of the  $k$  nearest neighbors is set to be the current prototype. The algorithm associates this prototype with the original candidate points. We repeat this step until the prototype moves less than a threshold  $\varepsilon_1$  or whenever the algorithm have reached  $j_{\max}$  iterations.

The complexity for the exact nearest neighbors search of a single point is  $n \log(n)$ . Applied to every data point multiple times, this complexity increases to  $n^2 j_{\max} \log(n)$ , preventing its application on Big datasets.

## 4 Model proposition

### 4.1 Approximate nearest neighbors search for density gradient ascent

One promising algorithmic complexity reduction approach relies on computing approximate nearest neighbors rather than exact neighbors. Among the techniques that can be used, Locality Sensitive Hashing (LSH), introduced in [4, 5], is a probabilistic method based on a random scalar projection of multivariate data point  $\mathbf{x}$  defined below:

$$L(\mathbf{x}; w) = (\mathbf{Z}^T \mathbf{x} + U)/w$$

where  $\mathbf{Z} \sim N(0, \mathbf{I}_d)$  is a standard  $d$ -variate normal random variable and  $U \sim \text{Unif}(0, w)$  is a uniform random variable on  $[0, w)$ ,  $w > 0$ . The LSH is parametrized by the number of buckets  $M_1$  in the hash table. In our context, we propose to set  $w = 1$ , and without loss of generality  $L_i \equiv L(\mathbf{X}_i; 1)$ . These scalar projections are sorted into their order statistics  $L_{(1)} < \dots < L_{(n)}$ , and their range is divided into  $M_1$  partition intervals of width  $w = (L_{(n)} - L_{(1)})/M_1$  where  $I_j = [L_{(1)} + w(j-1), L_{(1)} + wj]$ ,  $\forall j \in \{j = 1, \dots, M_1\}$ . The hash value of  $\mathbf{x}$  is the index of the interval in which  $L(\mathbf{x}; 1)$  falls

$$H(\mathbf{x}) = j \mathbf{1}\{L(\mathbf{x}; 1) \in I_j\},$$

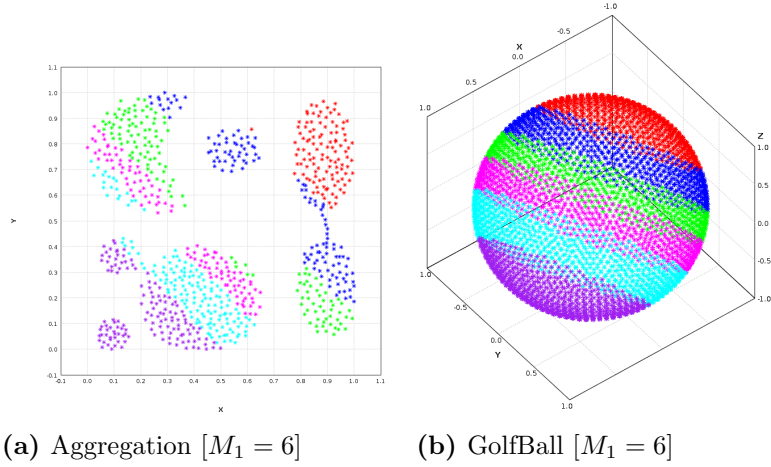
where  $\mathbf{1}\{\cdot\}$  is the indicator function. To search for approximate nearest neighbors, the reservoir of potential nearest neighbors is set to the bucket which contains the hash value. This reservoir is enlarged if necessary by concatenating the adjacent buckets. The approximate  $k$  nearest neighbors of  $\mathbf{x}$  are the  $k$  nearest neighbors only drawn from the reduced reservoir  $R(\mathbf{x})$  defined as below:

$$k\text{-}\widetilde{\text{nn}}(\mathbf{x}) = \{\mathbf{X}_i \in R(\mathbf{x}) : \|\mathbf{x} - \mathbf{X}_i\| \leq \delta_{(k)}(\mathbf{x})\},$$

where  $\delta_{(k)}(\mathbf{x})$  is the  $k$ th nearest neighbor distance to  $\mathbf{x}$ . The approximation error in the nearest neighbors to  $\mathbf{x}$  induced by searching in  $R(\mathbf{x})$  rather than the full dataset is probabilistically controlled [6]. Our work introduced an improvement to the classical LSH, based on the following observations, as follows:

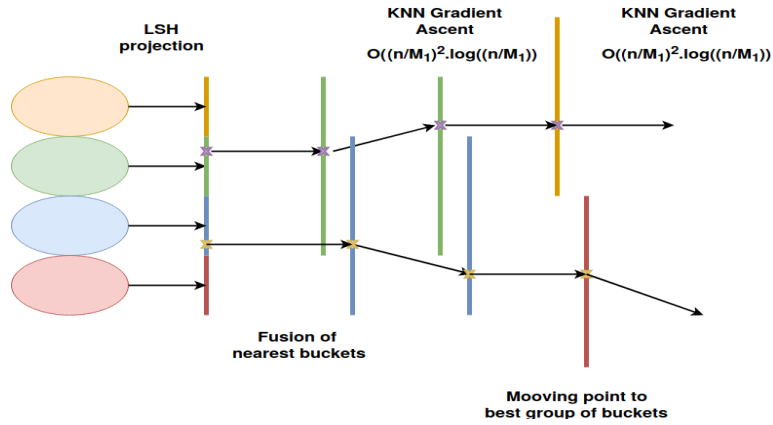
- it takes into account the properties of the local data space more accurately. Rather than looking exclusively in the bucket where the prototype lies, we also look in the neighbor buckets on both sides (i.e.  $p = 1$ ). The main advantage of this is to take into account the case where a prototype is at the border of a bucket and so some of its  $k$  nearest neighbors mostly likely fall into the neighbor buckets. This is especially important for high values of  $k$ . It is computationally more expensive but the cost can be controlled for a fixed bucket size. The memory cost is increased by a factor of  $2p + 1$  per partition due to copying the neighboring layers into the active one.

The LSH method partitions the data space into buckets of approximately  $k$  nearest neighbors, which are delimited by parallel hyperplanes. In practice, the LSH controls the number of neighboring buckets to two, except for the edge buckets which have only one neighbor bucket. This is in contrast to cell based buckets, where the number of neighbor buckets increases exponentially with the number of dimensions. Fig. 1 illustrates the LSH buckets of approximate nearest neighbors on 2D and 3D data examples: the orientation of hyperplanes depends on the random projections utilized to construct the buckets.



**Figure 1:** LSH buckets for the Aggregation and GolfBall datasets. The label  $[M_1]$  indicates the LSH with  $M_1$  buckets.

- One add the possibility of allowing the prototype to change buckets during its gradient ascent. In this case, we look for its  $k_2$  nearest neighbors in order to place the prototype within the most representative bucket within a majority voting process. Thus a prototype can pass through multiple buckets before converging to its final position, as illustrated in Fig. 2.



**Figure 2:** Passage of a prototype through different LSH buckets during the gradient ascent.  $n$  and  $M_1$  indicate respectively the number of observation and the number of buckets

Algorithm 3 describes the NNGA<sup>+</sup>, an approximate nearest neighbor search using LSH with the hash function  $H$ . The inputs are the data sample  $\mathbf{X}_1, \dots, \mathbf{X}_n$ , the candidate points  $\mathbf{x}_1, \dots, \mathbf{x}_m$ , and the tuning parameters: the number of nearest neighbors  $k_1$  and the number of buckets in the hash

table  $M_1$ . In line 1, the hash table is created by applying the LSH to the data values  $\mathbf{X}_1, \dots, \mathbf{X}_n$ . In lines 2–6, for each candidate point  $\mathbf{x}_\ell$ , the approximate  $k_1$  nearest neighbors  $k\text{-}\widetilde{\text{nn}}(\mathbf{x}_\ell)$  are computed from within the reservoir  $R(\mathbf{x}_\ell)$ .

---

**Algorithm 3** NNGA<sup>+</sup> – Approximate Nearest Neighbors Gradient Ascent with LSH and adjacent buckets

---

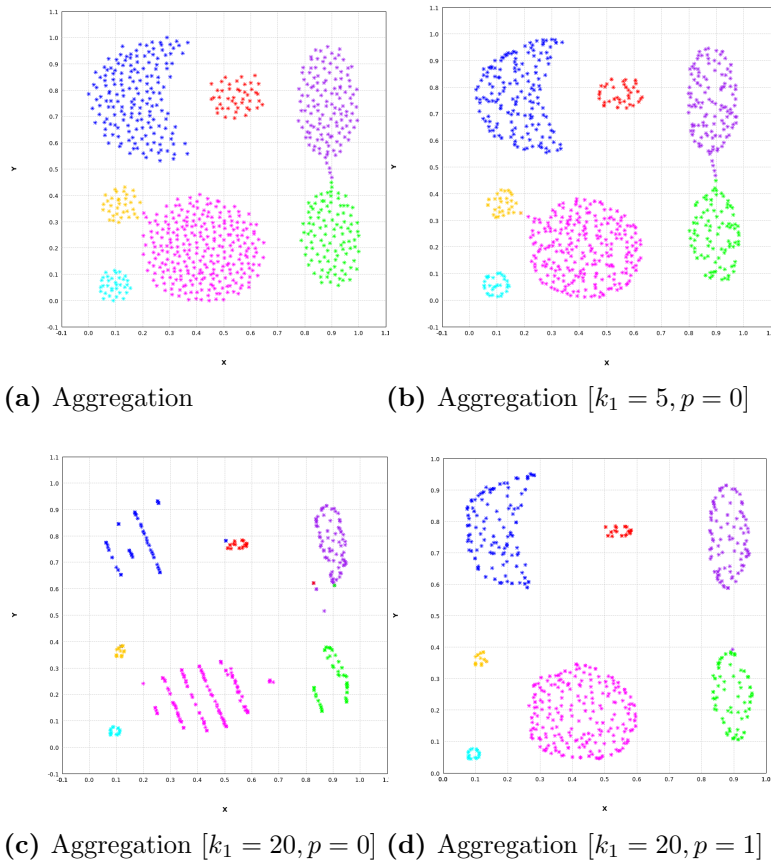
**Input:**  $\{\mathbf{X}_1, \dots, \mathbf{X}_n\}, \{\mathbf{x}_1, \dots, \mathbf{x}_m\}, k_1, M_1$   
**Output:**  $\{k\text{-}\widetilde{\text{nn}}(\mathbf{x}_1), \dots, k\text{-}\widetilde{\text{nn}}(\mathbf{x}_m)\}$   
/\* Create hash table with  $M_1$  buckets \*/  
1: **for**  $i := 1$  to  $n$  **do**  $H_i := H(\mathbf{X}_i)$ ;  
/\* Search for approx  $nn$  in adjacent buckets \*/  
2: **for**  $\ell := 1$  to  $m$  **do**  
3:      $R(\mathbf{x}_\ell) := \{\mathbf{X}_i : H_i = H(\mathbf{x}_\ell), i \in \{1, \dots, n\}\}$   
4:     **while**  $\text{card}(R(\mathbf{x}_\ell)) < k_1$  **do**  
5:          $R(\mathbf{x}_\ell) := R(\mathbf{x}_\ell) \cup \text{neighbor bucket}$ ;  
6:      $k\text{-}\widetilde{\text{nn}}(\mathbf{x}_\ell) := k\text{-nn}$  from  $R(\mathbf{x}_\ell)$  to  $\mathbf{x}_\ell$ ;

---

Fig.3 illustrates the effect of including neighbor buckets ( $p = 1$ ) or not ( $p = 0$ ) in the NNGA<sup>+</sup>. Fig. 3a demonstrates the effective of our algorithm on the Aggregation dataset. In Figure 3b a low value of  $k_1$  is used and no neighbor buckets are allowed where the data shrinks slightly. If we increase the value of  $k_1$  to 20 with no neighbor buckets, in Figure 3c, then we observe that the data are artificially forced to follow the hyperplanes which delimit the different buckets. Figure 3d shows our algorithmic improvement with adding one layer of neighbor buckets where the underlying structure of data is maintained.

It is important to note that the use of the LSH to reduce the complexity of kernel Mean Shift clustering was already proposed in [17], but authors did not quantify the reduction in complexity. The complexity of our NNGA<sup>+</sup> is reduced to  $O((\frac{n}{M_1})^2 \log(\frac{n}{M_1}))$  per bucket with  $M_1$  buckets, and so the total complexity is  $O((\frac{n}{M_1})^2 \log(\frac{n}{M_1}))$  for all buckets. Because of this segmentation of the original data space into  $M_1$  sub-spaces, the complexity is inversely proportional to the number of buckets. The trade-off is that the data points in each bucket have to be sufficiently representative of the local properties of the original space. Thus the number of buckets  $M_1$  is a crucial tuning parameter. Despite this, there are no optimal methods for selecting the number of buckets [18]. Consequently, we will examine empirical choices of the number of buckets to study the performance of our proposed method in section 5.





**Figure 3:** Results of NNGA<sup>+</sup> for the Aggregation dataset.  $[k_1, p]$  indicates NNGA<sup>+</sup> with  $k_1$  nearest neighbors  $k_1$  and  $p$  neighbor bucket layers.

## 4.2 Cluster labeling: $\varepsilon$ -proximity

NNGA<sup>+</sup> carries the gradient ascent on the data points until they have converged to their prototypes. The question is then how to label the data point with clusters. A first solution is to assign the same cluster to all points sharing the same prototype. As observed in Figure 3d, even with an a priori good choice of  $k_1$ , there are (possibly) hundreds of generated prototypes, so assigning a label to each point according to its closest prototype is not effective because of generating too many clusters. Applying the density gradient ascent NNGA<sup>+</sup> leads to a converged dataset with increased inter-cluster distances and decreased intra-cluster distances as compared to the original dataset. In order to further exploit this property, we propose a new proximity-based approach where points which are under a threshold  $\varepsilon$  from each other are considered to belong to the same cluster.

The aim of  $\varepsilon$ -proximity algorithm is to gather all points which are under a distance  $\varepsilon$  from each other. In order to apply this method, we have to

build the similarity matrix which has a  $O(n^2)$  time complexity, preventing any Big Data application. For a scalable approach, we apply the LSH on the dataset and generate clusters on these LSH buckets through the local version of  $\varepsilon$ -proximity clustering.

Algorithm 4 illustrates the local version of the algorithm. It consists in exploring the similarity matrix  $\mathbf{S}$  which is defined as a map whose objects IDs are the keys and whose pairs (object IDs, distances) are the values. We initialize the process by taking the first object of  $\mathbf{S}$  and cluster with it every point whose distance is less than  $\varepsilon_2$ . We then apply this exploration process by iteratively adding the  $\varepsilon_2$  nearest neighbors of these added points until this process terminates. During the process we remove the explored points from  $\mathbf{S}$  to avoid repeated calculations. Once the first cluster is generated, we take another object from outside this first cluster from the reduced similarity matrix  $\mathbf{S}$  and repeat the above cluster formation, until all objects are assigned into a cluster label.

The distributed version consists to apply this algorithm in each LSH partitions, we merge with each bucket its right or left neighbor bucket, maintaining the bucket order. Once this step is completed, we apply a *MapPartitions* procedure where we check if two clusters of two different buckets share  $k_3$  pairs of points under  $\varepsilon$  (we use  $k_3 = 1$  throughout this paper). Thus these two clusters are considered to form a single cluster. We obtain a dataset which chains common clusters between partitions: all chained clusters are assigned with the same label by generating an undirected graph where each connected subgraph represents a cluster. The search for connected components in a graph is a common problem which can be solved in  $O(n)$   $n$  being the number of vertices.

It is important to bound the buckets size because the local version of  $\varepsilon$ -proximity clustering and the check for cluster fusion between two buckets have quadratic complexity in this size. Empirically we advise to set the number of buckets in order to have around 500 to 2000 data points in each bucket.

A notable problem still remains with the choice of the main tuning parameter  $\varepsilon$ . We set it to be the average of distance from each point to their  $k$  nearest neighbors. We compute it as an approximate value in using LSH procedure in order to maintain the scalability property.

## 5 Numerical experiments

Our experiments are carried out on the Grid'5000 testbed which is the french national testbed for computer science research. It allows the deployment of a user-specified operating system within the Grid'5000 hardware. We use a dedicated Spark Linux image optimized for Grid'5000 where Apache Spark is deployed on top of Spark in Standalone mode. Apache Spark is a fast general

---

**Algorithm 4** Local  $\varepsilon$ -proximity labeling

---

```
Input:  $\{\mathbf{x}_1, \dots, \mathbf{x}_m\}, \mathbf{S}, \varepsilon_2$   
Output:  $\{\tilde{c}(\mathbf{x}_1), \dots, \tilde{c}(\mathbf{x}_m)\}$   
1: needToVisit  $\leftarrow$  Set(S.head)  
2:  $c_{ID} \leftarrow 0$   
3: clusters  $\leftarrow$  Map.empty[Int, Set[Int]]  
4: clusters += ( $c_{ID}$ , needToVisit)  
5: while S has elements do  
  /*  $p_c$  is the current point of needToVisit */  
6:   for each  $p_c$  in needToVisit do  
  /* Add all points under  $\varepsilon_2$  to needToVisit */  
7:     needToVisitUpdated  $\leftarrow$   $\{p \in \mathbf{S}, \text{dist}(p, p_c) \leq \varepsilon_2\}$   
  /* Remove explored point from the similarity matrix */  
8:     S -=  $p_c$   
  /* Update points to explore on next iterations */  
9:     needToVisit  $\leftarrow$  needToVisitUpdated  
10:    if needToVisit is empty then  
11:       $c_{ID} += 1$   
12:      needToVisit  $\leftarrow$  Set(S.head)  
  /* Create a new entry in the clusters map */  
13:      clusters += ( $c_{ID}$ , needToVisit)  
14:    else  
  /* Add new points to cluster  $c_{ID}$  */  
15:      clusters( $c_{ID}$ ) += needToVisit
```

---

purpose distributed computing system based on a master-slaves architecture. Only the deployment of the image is automatized. We manually reserve the nodes and provide the Spark cluster with our code to execute the different experiments on a  $2 \times 8$  core Intel Xeon E5-2630v3 CPUs and 128 Gb RAM set up. We repeat each experiment ten times for robustness.

A key concept in Spark is the resilient distributed dataset (RDD) which is a read-only collection of objects partitioned across a group of machines which can be rebuilt if necessary from the hierarchy of previous RDD operations. Most of the *Map* and *Reduce* operations will be performed on RDDs even if other pure Scala *Map* and *Reduce* operations are executed inside each Spark partition. We implement our algorithm in Scala because it is the Spark's native language and thus allows for good performance. Most tuning parameters of our algorithm have an impact on both the execution time and the cluster labeling quality. We focus our study on the dataset size  $n$ , the maximum number of Gradient Ascent iterations  $j_{\max}$ , the number of LSH buckets  $M_1$ , the threshold  $\varepsilon_1$ , and the number of nearest neighbors  $k_1$ .

## 5.1 Datasets and tuning parameters

We use a range from 2 to high dimensional datasets with different sizes, as summarized in Table 1 [19, 20]. To ensure the comparability of the results across these different datasets, all algorithms are carried out on the normalized version of the datasets:  $x_i = (x_i - x_i^{\min}) / (x_i^{\max} - x_i^{\min})$  where  $x_i$  is the  $i$ th component of  $\mathbf{x}$ , and  $x_i^{\min}, x_i^{\max}$  are respectively the  $i$ th marginal minimum and maximum values. Table 2 shows the combination of tuning parameters used in the comparison between the gradient ascent (NNGA<sup>+</sup>),  $k$ -means [12], DBScan [13] and  $\varepsilon$ -proximity.

Dataset	$n$	$d$	$N$
R15	600	2	15
Aggregation	788	2	7
Sizes5	1000	2	4
EngyTime	4096	2	2
Banana	4811	2	2
S3	5000	2	15
Disk6000	6000	2	2
DS1	9153	2	14
Hepta	212	3	7
Hyperplane	100000	10	5
CovType10	581012	10	7
image <sup>1</sup>	154401	5	–
Own image	5000000	5	–
ScalabiltyDS	140000000	10	–

**Table 1:** Experimental datasets.  $n$  is the dataset size,  $d$  is the data dimension,  $N$  is the number of clusters.

## 5.2 Evaluation of clustering quality

### 5.2.1 Quantitative evaluation

NNGA<sup>+</sup> is compared to the following cluster labeling techniques:  $k$ -means [12], DBScan [13], and  $\varepsilon$ -proximity. To evaluate the quality of the clustering, we use both the Normalized Mutual Information (NMI) [21] and the RAND index [22]. The value of each measure lies between 0 and 1. A higher value indicates better clustering results.

One notable observation of the impact of the NNGA<sup>+</sup> pre-processing step on clustering labeling concerns the datasets with Gaussian (i.e. ellipsoidal) clusters. As the results for the Hepta, R15, S3, and Sizes5 datasets as Table 3 and 4 show, applying NNGA<sup>+</sup> results in uniformly better clustering quality than without the NNGA<sup>+</sup>. NNGA<sup>+</sup> pre-processing in DBScan and

dataset	NNGA <sup>+</sup>	$k$ -means	DBScan	$\epsilon$ -proximity
Aggregation	without	$k = 7$	$\epsilon = 0.05,$ $minPts = 8$	$\epsilon_{knn} = 10,$ $M_1 = 8$
	with $k_1 = 40$			$\epsilon_{knn} = 30,$ $M_1 = 8$
Banana	without	$k = 2$	$\epsilon = 0.02,$ $minPts = 3$	$\epsilon_{knn} = 0.1,$ $M_1 = 8$
	with $k_1 = 40$			
Disk6000	without	$k = 7$	$\epsilon = 0.02,$ $minPts = 4$	$\epsilon_{knn} = 0.021,$ $M_1 = 8$
	with $k_1 = 40$			
DS1	without	$k = 14$	$\epsilon = 0.03,$ $minPts = 25$	$\epsilon_{knn} = 30,$ $M_1 = 8$
	with $k_1 = 40$			
EngyTime	without	$k = 2$	$\epsilon = 0.1,$ $minPts = 100$	$\epsilon_{knn} = 50,$ $M_1 = 8$
	with $k_1 = 40$			
Hepta	without	$k = 7$	$\epsilon = 0.1,$ $minPts = 10$	$\epsilon_{knn} = 10,$ $M_1 = 4$
	with $k_1 = 40$			$\epsilon_{knn} = 20,$ $M_1 = 4$
R15	without	$k = 15$	$\epsilon = 0.05,$ $minPts = 25$	$\epsilon_{knn} = 5,$ $M_1 = 8$
	with $k_1 = 40$			
S3	without	$k = 15$	$\epsilon = 0.05,$ $minPts = 50$	$\epsilon_{knn} = 15,$ $M_1 = 8$
	with $k_1 = 40$			$\epsilon_{knn} = 15,$ $M_1 = 8$
Sizes5	without	$k = 4$	$\epsilon = 0.08,$ $minPts = 8$	$\epsilon_{knn} = 10,$ $M_1 = 8$
	with $k_1 = 40$			$\epsilon_{knn} = 35,$ $M_1 = 8$
Unbalanced	without	$k = 8$	$\epsilon = 0.05,$ $minPts = 20$	$\epsilon_{knn} = 50,$ $M_1 = 8$
	with $k_1 = 40$		$\epsilon = 0.1,$ $minPts = 20$	$\epsilon_{knn} = 80,$ $M_1 = 8$
Hyperplane	without	$k = 5$	$\epsilon = 0.05,$ $minPts = 8$	$\epsilon_{knn} = 5,$ $M_1 = 100$
	with $k_1 = 40$			
CovType10	without	$k = 5$	$\epsilon = 0.05,$ $minPts = 8$	$\epsilon_{knn} = 20,$ $M_1 = 500$

**Table 2:** Clustering algorithms parameters.  $\epsilon = v_1$  stands for manual setting of  $\epsilon$  and  $\epsilon_{knn} = v_2$  defines the number of the  $k$  nearest neighbor for the  $\epsilon$  approximation.  $M_1 =$  represents the number of buckets during the LSH phase.

$\varepsilon$ -proximity clustering also leads to better (or at least as good) clustering accuracy than without NNGA<sup>+</sup> except for Disk6000 which has nested clusters. In this case, a high  $k_1$  value leads to the creation of links between clusters which are well separated. Finally we observe for the high dimensional datasets, Hyperplan and CovType10,  $\varepsilon$ -proximity with NNGA<sup>+</sup> pre-processing performs as well as the classical  $k$ -means, and it out-performs the DBScan which is unable to effectively cluster such huge datasets.

Data	NNGA <sup>+</sup>	$k$ -means	DBScan	$\varepsilon$ -proximity
Aggregation	without	0.83 ± 0.022	0.98 ± 0.00	0.89 ± 0.00
	with [ $k_1 = 50$ ]	<b>0.87 ± 0.06</b>	<b>0.99 ± 0.00</b>	<b>0.97 ± 0.02</b>
Banana	without	0.31 ± 0.00	<b>1.00 ± 0.00</b>	<b>1.00 ± 0.00</b>
	with [ $k = 40$ ]	<b>0.32 ± 0.01</b>	<b>1.00 ± 0.00</b>	<b>1.00 ± 0.00</b>
Disk6000	without	0.00 ± 0.00	<b>1.00 ± 0.00</b>	<b>1.00 ± 0.00</b>
	with [ $k = 100$ ]	<b>0.02 ± 0.00</b>	0.25 ± 0.00	0.32 ± 0.00
DS1	without	0.75 ± 0.01	<b>0.94 ± 0.00</b>	<b>0.97 ± 0.00</b>
	with [ $k_1 = 50$ ]	<b>0.77 ± 0.03</b>	<b>0.94 ± 0.00</b>	0.95 ± 0.00
EngyTime	without	<b>0.98 ± 0.00</b>	0.75 ± 0.00	0.01 ± 0.00
	with [ $k_1 = 200$ ]	0.97 ± 0.00	<b>0.97 ± 0.00</b>	<b>0.85 ± 0.08</b>
Hepta	without	0.98 ± 0.03	0.83 ± 0.00	<b>0.99 ± 0.02</b>
	with [ $k_1 = 20$ ]	<b>0.99 ± 0.02</b>	<b>1.00 ± 0.00</b>	0.97 ± 0.04
R15	without	0.96 ± 0.02	0.99 ± 0.00	<b>0.93 ± 0.01</b>
	with [ $k_1 = 20$ ]	<b>0.99 ± 0.00</b>	<b>0.99 ± 0.00</b>	0.91 ± 0.02
S3	without	0.78 ± 0.01	0.42 ± 0.00	0.09 ± 0.00
	with [ $k_1 = 40$ ]	<b>0.79 ± 0.00</b>	<b>0.75 ± 0.00</b>	<b>0.74 ± 0.00</b>
Sizes5	without	0.81 ± 0.12	0.80 ± 0.00	0.60 ± 0.12
	with [ $k_1 = 20$ ]	<b>0.89 ± 0.08</b>	<b>0.91 ± 0.00</b>	<b>0.89 ± 0.01</b>
Unbalanced	without	<b>0.94 ± 0.05</b>	<b>0.99 ± 0.00</b>	0.97 ± 0.01
	with [ $k_1 = 40$ ]	0.92 ± 0.05	<b>0.99 ± 0.00</b>	<b>0.98 ± 0.01</b>
Hyperplane	without	<b>0.01 ± 0.00</b>	One cluster	0.02 ± 0.00
	with [ $k_1 = 50$ ]	0.00 ± 0.00	found	<b>0.04 ± 0.00</b>
CovType10	without	0.07 ± 0.006	dataset is	0.03 ± 0.00
	with [ $k_1 = 50$ ]	<b>0.07 ± 0.02</b>	too massive	<b>0.09 ± 0.01</b>

**Table 3:** NMI clustering quality indices for the cluster labeling on the experimental datasets with and without prior application of NNGA<sup>+</sup>. The bold entries indicate the optimal dataset, and  $\pm$  entries are the standard deviations over 10 trials.

As shown in Figure. 4, our version of the nearest neighbors gradient ascent NNGA<sup>+</sup> results in shrinking the data points toward their local modes. The notation ‘with NNGA<sup>+</sup> [ $k_1$ ]’ indicates that we carried out NNGA<sup>+</sup>  $k_1$  nearest neighbors in the Mean Shift gradient ascent. DS1 with NNGA<sup>+</sup> [ $k_1 = 50$ ] presents a more compact version than its original version, maintaining the underlying data structures whilst increasing empty space between clusters. Likewise for the Hepta dataset. One notable difference arises for the uniform clusters in Disk6000. In the latter, we observe formation of empty space next to high density areas, whereas in the former, the data points gather towards the central modes.

An application of NNGA<sup>+</sup> could result in one of two major cases depending on the nature of clusters. If the points are homogeneously distributed over a cluster, a higher value of  $k_1$  results in more empty space surrounding the high density regions. On the other hand, if the points are Gaussian distributed, they will converge efficiently towards the high density region,

Data	NNGA <sup>+</sup>	$k$ -means	DBScan	$\varepsilon$ -proximity
Aggregation	without	0.91 $\pm$ 0.01	0.99 $\pm$ 0.00	0.93 $\pm$ 0.00
	with [ $k_1 = 50$ ]	<b>0.93 <math>\pm</math> 0.04</b>	<b>1.00 <math>\pm</math> 0.00</b>	<b>0.98 <math>\pm</math> 0.02</b>
Banana	without	0.70 $\pm$ 0.00	<b>1.00 <math>\pm</math> 0.00</b>	<b>1.00 <math>\pm</math> 0.00</b>
	with [ $k = 40$ ]	<b>0.70 <math>\pm</math> 0.01</b>	1.00 $\pm$ 0.00	1.00 $\pm$ 0.00
Disk6000	without	<b>0.50 <math>\pm</math> 0.00</b>	<b>1.00 <math>\pm</math> 0.00</b>	<b>1.00 <math>\pm</math> 0.00</b>
	with [ $k = 100$ ]	0.50 $\pm$ 0.00	0.33 $\pm$ 0.00	0.34 $\pm$ 0.00
DS1	without	0.86 $\pm$ 0.00	0.98 $\pm$ 0.00	0.98 $\pm$ 0.00
	with [ $k_1 = 50$ ]	<b>0.87 <math>\pm</math> 0.01</b>	<b>0.98 <math>\pm</math> 0.00</b>	<b>1.00 <math>\pm</math> 0.00</b>
EngyTime	without	<b>1.00 <math>\pm</math> 0.0</b>	0.90 $\pm$ 0.00	0.50 $\pm$ 0.00
	with [ $k_1 = 200$ ]	0.99 $\pm$ 0.00	<b>0.99 <math>\pm</math> 0.00</b>	<b>0.93 <math>\pm</math> 0.05</b>
Hepta	without	0.99 $\pm$ 0.02	0.94 $\pm$ 0.00	<b>0.99 <math>\pm</math> 0.02</b>
	with [ $k_1 = 20$ ]	<b>0.99 <math>\pm</math> 0.02</b>	<b>1.00 <math>\pm</math> 0.00</b>	0.98 $\pm$ 0.04
R15	without	0.99 $\pm$ 0.01	1.00 $\pm$ 0.00	<b>0.99 <math>\pm</math> 0.00</b>
	with [ $k_1 = 20$ ]	<b>1.00 <math>\pm</math> 0.00</b>	<b>1.00 <math>\pm</math> 0.00</b>	0.98 $\pm$ 0.00
S3	without	0.96 $\pm$ 0.00	0.52 $\pm$ 0.00	0.12 $\pm$ 0.00
	with [ $k_1 = 40$ ]	<b>0.96 <math>\pm</math> 0.00</b>	<b>0.95 <math>\pm</math> 0.00</b>	<b>0.96 <math>\pm</math> 0.00</b>
Sizes5	without	0.88 $\pm$ 0.13	0.97 $\pm$ 0.00	0.83 $\pm$ 0.07
	with [ $k_1 = 20$ ]	<b>0.95 <math>\pm</math> 0.09</b>	<b>0.98 <math>\pm</math> 0.00</b>	<b>0.97 <math>\pm</math> 0.00</b>
Unbalanced	without	<b>0.97 <math>\pm</math> 0.03</b>	<b>1.00 <math>\pm</math> 0.00</b>	<b>1.00 <math>\pm</math> 0.00</b>
	with [ $k_1 = 40$ ]	0.96 $\pm$ 0.03	<b>1.00 <math>\pm</math> 0.00</b>	<b>1.00 <math>\pm</math> 0.00</b>
Hyperplane	without	<b>0.62 <math>\pm</math> 0.00</b>	One cluster	0.30 $\pm$ 1.00
	with [ $k_1 = 50$ ]	0.62 $\pm$ 0.00	found	<b>0.31 <math>\pm</math> 0.00</b>
CovType10	without	<b>0.59 <math>\pm</math> 0.00</b>	dataset is	0.38 $\pm$ 0.00
	with [ $k_1 = 50$ ]	0.58 $\pm$ 0.01	too massive	<b>0.56 <math>\pm</math> 0.04</b>

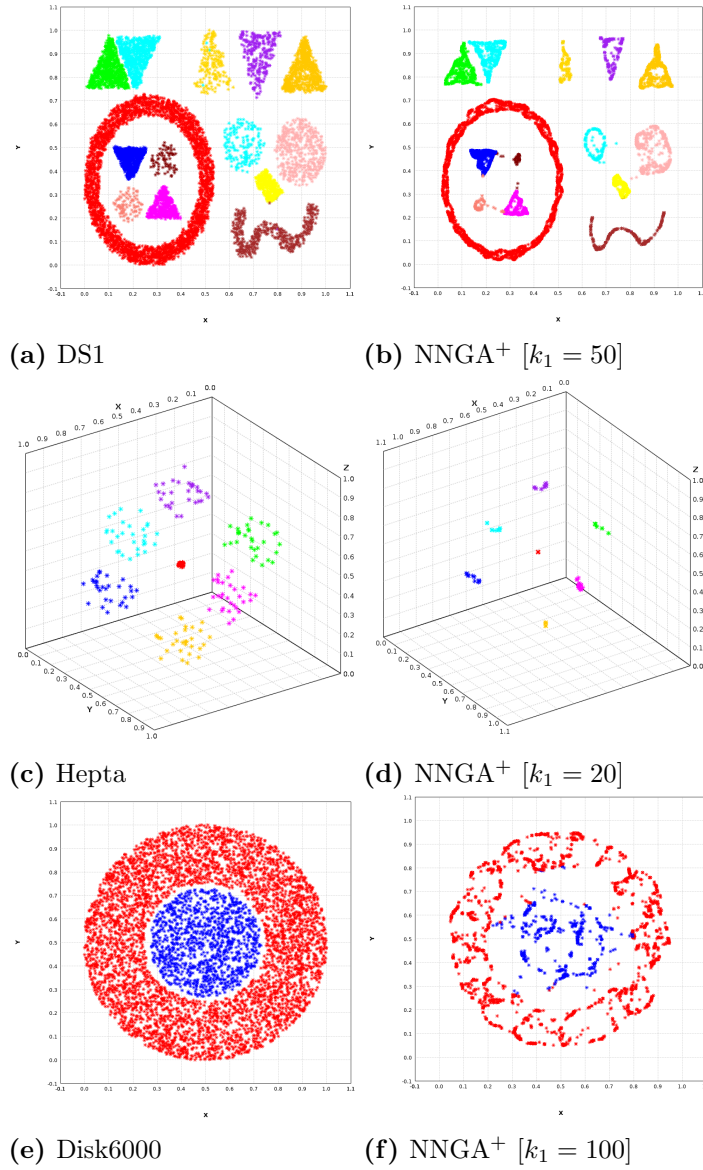
**Table 4:** RAND clustering quality indices for the cluster labeling on the experimental datasets with and without prior application of NNGA<sup>+</sup>. The bold entries indicate the optimal dataset, and  $\pm$  entries indicate the standard deviations over 10 trials.

reducing the noise whilst increasing the inter-cluster distances.

Figure 5 presents some results obtained by applying  $k$ -means on the original data points and converged data points obtained after applying the NNGA<sup>+</sup> algorithm. Figure 5a shows an application of  $k$ -means on Sizes5 dataset with  $k = 4$  clusters without NNGA<sup>+</sup>. Even if the number of clusters in the  $k$ -means is the correct number, these four clusters do not match so closely the original clusters (NMI=0.81, RAND=0.88). Figure 5b shows the  $k$ -means cluster labeling results applied on NNGA<sup>+</sup> output with  $k_1 = 20$ . We observe visually that these clusters are more similar to the original clusters (NMI=0.89, RAND=0.95).

DBScan cluster labeling collates points more efficiently after NNGA<sup>+</sup> is applied than without NNGA<sup>+</sup>, as shown in Figure 6. NNGA<sup>+</sup> is an efficient way to attach noisy points to their closest cluster. Furthermore, NNGA<sup>+</sup> facilitates more robust choices for the DBScan parameters, since it increases the local density which improves the detection of smaller clusters.

The results for  $\varepsilon$ -proximity labeling are presented in Figure 7. For the Aggregation dataset, the NNGA<sup>+</sup> improves the clustering by being able to distinguish between clusters which are joined a filamentary structure. Without NNGA<sup>+</sup> these filamentary structures cause problems. An analogous improvement in clustering quality is also observed for the R15 dataset. The cluster labelling is similar for the Unbalance dataset with or without NNGA<sup>+</sup>.



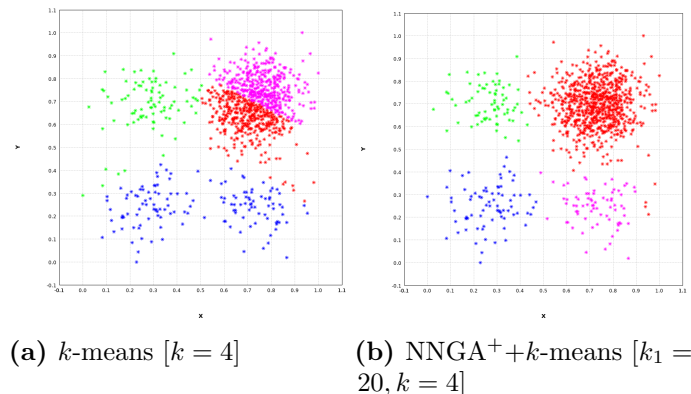
**Figure 4:** Gradient ascent NNGA<sup>+</sup> on the DS1, Hepta and Disk6000 datasets. The label ‘NNGA<sup>+</sup> [ $k_1$ ]’ indicates the application of NNGA<sup>+</sup> with  $k_1$  nearest neighbors.

On the other hand, for the DS1 dataset, the prior application of NNGA<sup>+</sup> leads to a decrease in the cluster labeling accuracy.

### 5.2.2 Visual evaluation with image segmentation

A resurgence in interest in the variant of the mean shift algorithm is due to its application to image segmentation [23] where an image is transformed into a color space in which clusters correspond to segmented regions in the original





**Figure 5:**  $k$ -means cluster labeling on Sizes5 datasets with and without prior application of NNGA<sup>+</sup>. The label ‘ $k$ -means [ $k$ ]’ indicates  $k$ -means cluster labeling with  $k$  clusters without NNGA<sup>+</sup>, and ‘NNGA<sup>+</sup>+ $k$ -means [ $k_1, k$ ]’ indicates NNGA<sup>+</sup> with  $k_1$  nearest neighbors followed by  $k$ -means cluster labeling with  $k$  clusters.


image. The 3-dimensional  $L^*u^*v^*$  color space [24, Eqs. 3.5-8a–f] is a common choice. Since an image is a 2-dimensional array of pixels, let  $(x, y)$  be the row and column index of a pixel. The spatial and color (range) information of a pixel can be concatenated into a 5-dimensional vector  $(x, y, L^*, u^*, v^*)$  in the joint spatial-range domain. An image segmentation algorithm based on the kernel mean shift was introduced in [25] which we adapt for use with NNGA<sup>+</sup>.

Our test image is image #36 from the colour training set from the Berkeley Segmentation Dataset and Benchmark<sup>2</sup>. Figure 8 shows the original RGB  $481 \times 321$  pixels JPEG image. The tuning parameters for the NNGA<sup>+</sup> are  $k_1 = 60$ ,  $j_{\max} = 15$ . We compute the NNGA<sup>+</sup>- $M_1$  with  $M_1 = 200, 400, 1000$  buckets.

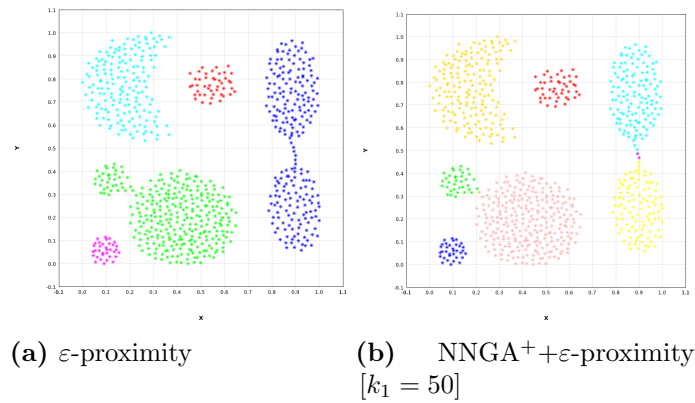
For the NNGA<sup>+</sup>-200 and NNGA<sup>+</sup>-400 in Figures 9b and 9d where we approximate nearest neighbours with respectively  $M_1 = 200$  and  $M_1 = 400$  buckets, some finer details are visible, such as the podia. For NNGA-200 and NNGA-400 in Figures 9a and 9c, the green background color bleeds into the starfish arms. For NNGA-1000 in Figure 9e, the starfish is not visible anymore. However, by taking more neighbor layers (NNGA<sup>+</sup>-1000  $p = 2$ ) in Figure 9f the starfish shape being delimited from the background.

The Berkeley Segmentation Dataset and Benchmark provides human expert segmentations of their images for comparisons. Figure 10a,b depict two edge detections made by Users #1109 and #1119. User #1109 focuses on segmenting the shape of the starfish, whilst ignoring the detail of the podia, whereas User #1119 concentrates on segmenting the individual podia in the foreground. We focus on the NNGA<sup>+</sup>-200 (Figure 9b), which segmentation

<sup>2</sup><http://www.eecs.berkeley.edu/Research/Projects/CS/vision/bsds>



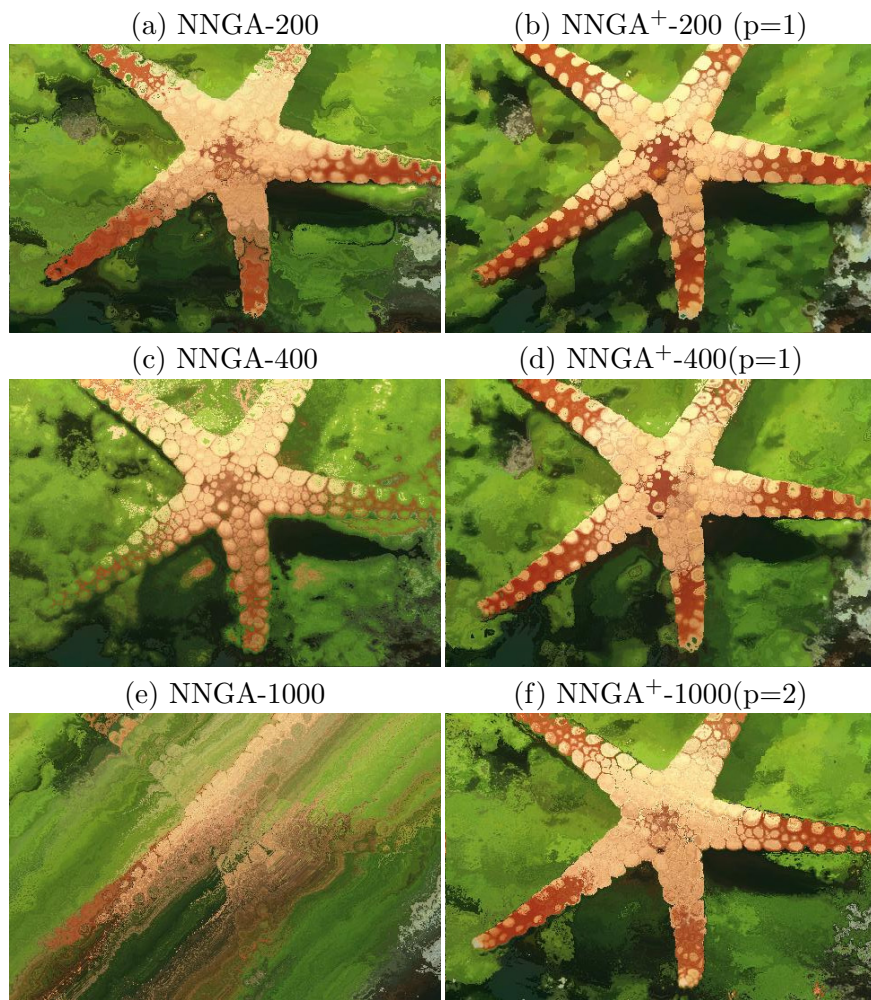
**Figure 6:** DBScan labeling on the EngyTime and Disk6000 datasets with and without prior application of  $\text{NNGA}^+$ . The label ‘DBScan’ indicates DBScan cluster labeling without  $\text{NNGA}^+$ , and ‘ $\text{NNGA}^+ + \text{DBScan } [k_1]$ ’ indicates  $\text{NNGA}^+$  with  $k_1$  nearest neighbors followed by DBScan cluster labeling.



**Figure 7:**  $\varepsilon$ -proximity cluster labeling on the Aggregation. The left image indicates  $\varepsilon$ -cluster labeling without  $\text{NNGA}^+$ . the right image ‘ $\text{NNGA}^+ + \varepsilon$ -proximity  $[k_1]$ ’ indicates  $\text{NNGA}^+$  with  $k_1$  nearest neighbors followed by  $\varepsilon$ -proximity labeling.

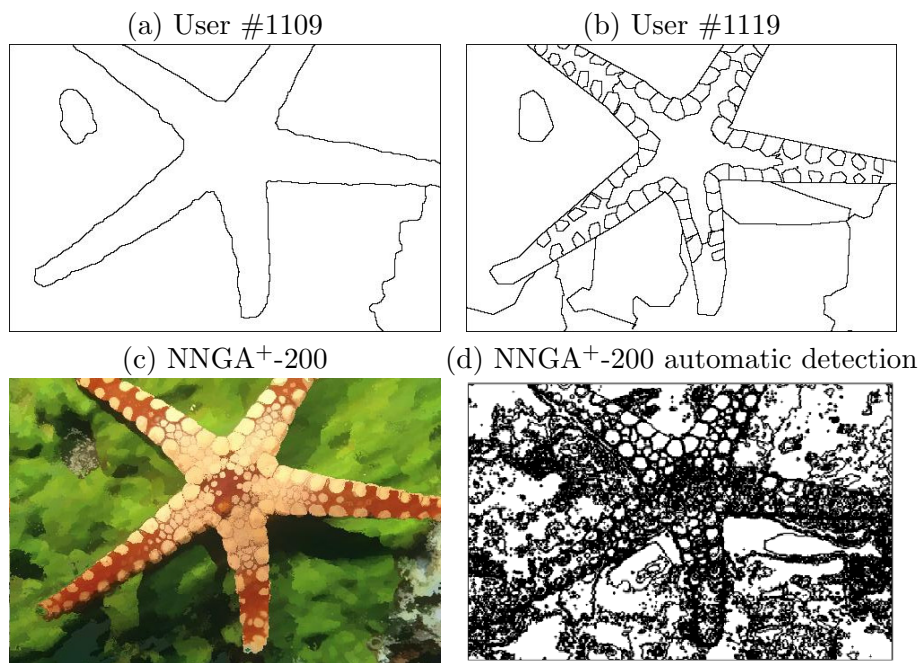


**Figure 8:** Original picture #36.



**Figure 9:** Colour image segmentation using in the left NNGA and in the right NNGA<sup>+</sup> with  $M_1 = 200, 400$  and  $1000$  buckets.

is closer to User #1119. Whilst this automatic edge detection (Figure 10d) in its current state remains too fragmented to be useful for a visual analysis, we anticipate that further application of statistical and image analyses will be able to improve it.

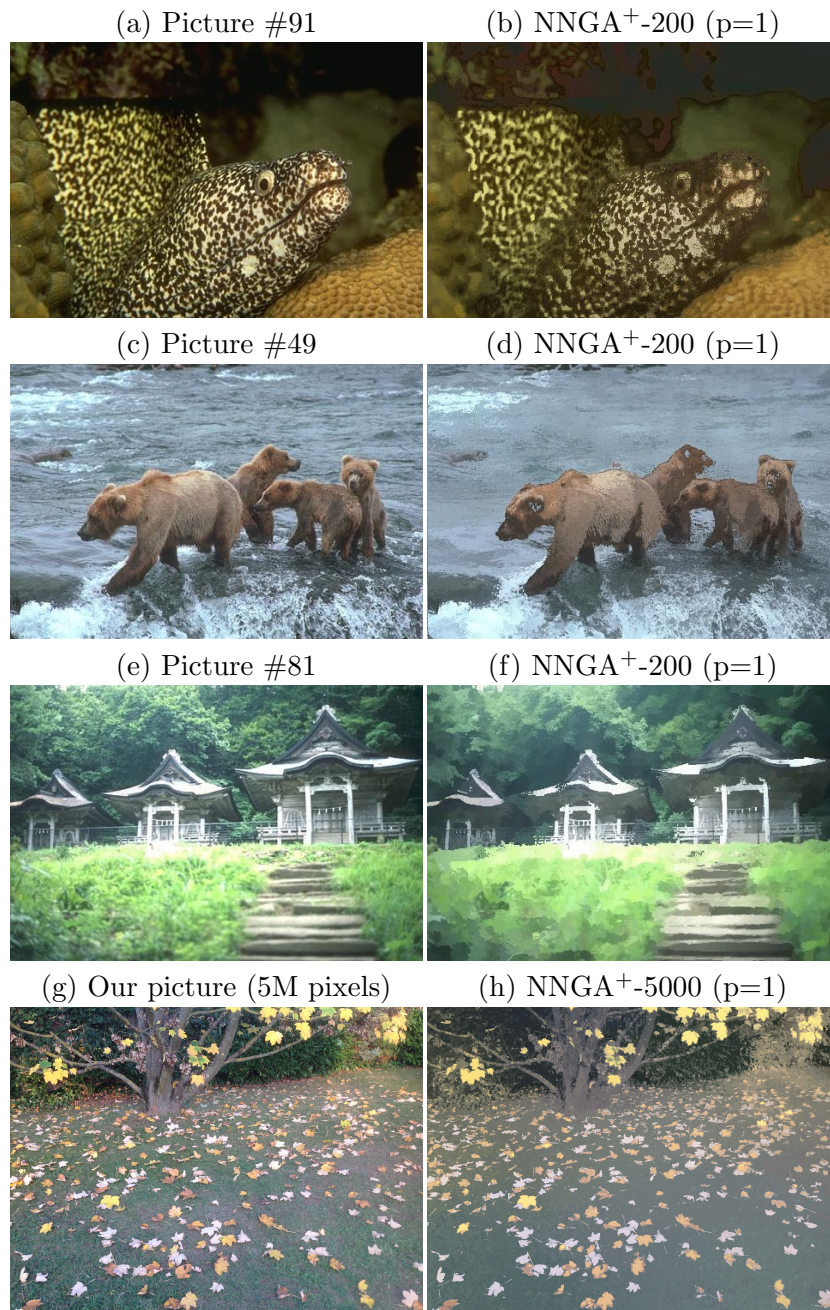


**Figure 10:** Edge detection of segmented images. (a,b) Two human experts: users #1109 and #1119. (c) NNGA<sup>+</sup>-200. (d) NNGA<sup>+</sup>-200 with automatic detection.

We observe from Figures 9 that  $M_1 = 200$  buckets is a suitable empirical choice and so we apply it to further images from the Berkeley image database as in Figure 11. We apply also the NNGA<sup>+</sup> on a picture of 5 millions pixels (taken by ourselves) compared to 150 thousand pixels of Berkeley’s pictures. We increase the number of buckets accordingly with the size of the picture to reach approximately 1000 data points per bucket. The image shows a good segmentation between the leaves and the rest of the picture. We distinguish the trunk from the background foliage too without difficulty showing the effectiveness of the algorithm even with big image.

### 5.3 Evaluation of scalability

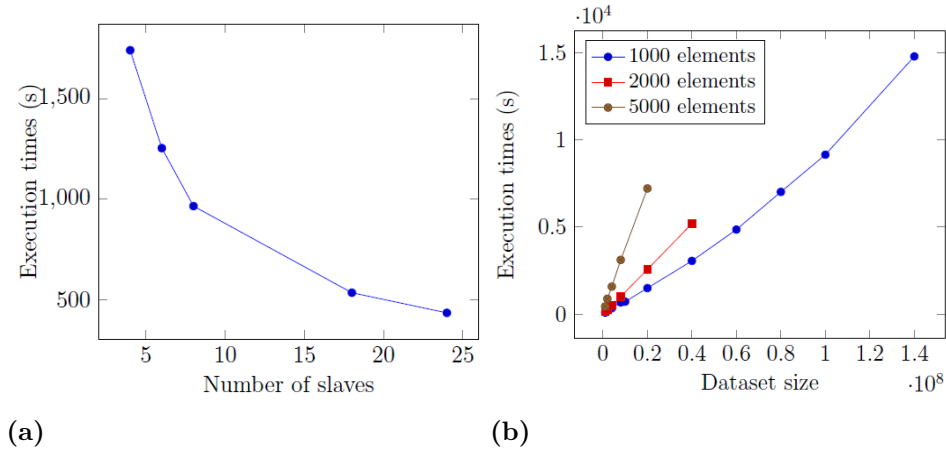
**Density gradient ascent** As the data are distributed over all buckets or Spark partitions, this allows for efficient computations in the local environment of a data point. Figure 12a shows that the execution time gradually decreases as the number of slaves increases for a datasets of a fixed size. In Figure 12b, for a fixed number of nodes, if we maintain the constant number



**Figure 11:** Further examples of segmented images with NNGA<sup>+</sup>-200.

of elements per bucket, the execution time grows linearly with the size of the dataset. This indicates that the number of nearest neighbor  $k_1$  needs to be constrained.

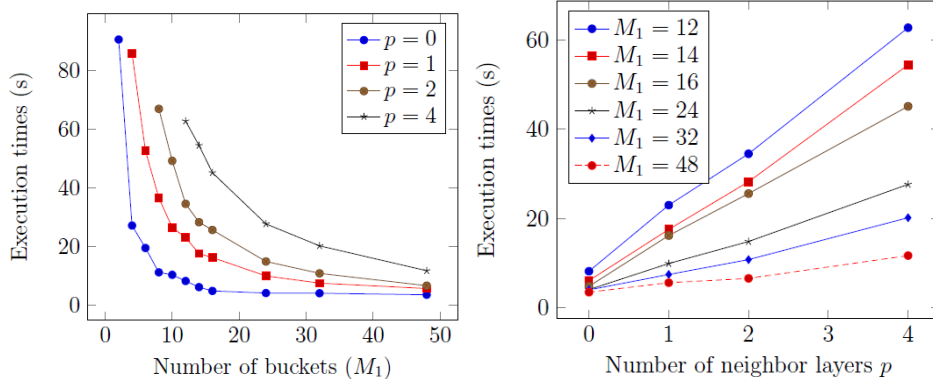
Since we have an  $O((\frac{n}{M_1})^2 \log(\frac{n}{M_1}))$  complexity per bucket for NNGA<sup>+</sup>, a suitable value for the number of buckets  $M_1$  is to keep the  $\frac{n}{M_1}$  ratio approximately equal to a constant  $C$ . Then the time complexity of NNGA<sup>+</sup> reduces to  $O(nC \log(C))$ . The scalability is demonstrated by the decrease in execution time with the number of slaves and a linear increase of execution time with the dataset size, reaching 140 million data points which is infeasible for the original quadratic algorithm.



**Figure 12:** (a) Execution times for NNGA<sup>+</sup> with respect to the number of slaves. (b) Execution times for NNGA<sup>+</sup> with respect to the dataset size  $n$  for fixed size buckets.

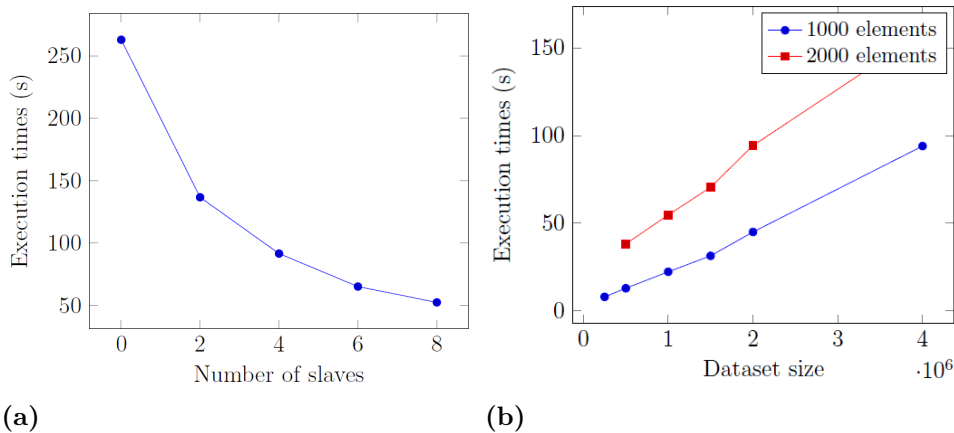
**LSH buckets and neighbor layers** For a fixed number of neighbor layers, we observe in Figure 13a that the execution time rapidly decreases and then slows down to reach a plateau, as the number of buckets increases. The observed plateau is due to the quadratic complexity of the NNGA<sup>+</sup>: more buckets leads to fewer data points within each bucket and so the execution times can quickly reach the minimal plateau after a sufficiently large number of buckets. We also studied the influence of the number of neighbors layers  $p$  on the execution time. Whilst NNGA<sup>+</sup> has quadratic time complexity in each bucket, if we select an appropriate number of nearest neighbors  $k_1$ , then we are able to control the execution time of NNGA<sup>+</sup> to be linear with respect to the number of neighbors layers  $p$ , as illustrated in Figure 13b.

**$\varepsilon$ -proximity cluster labeling** Concerning  $\varepsilon$ -proximity cluster labeling, similar remarks as for the gradient ascent apply here. As Fig.14 follows same decrease in the execution time as for the gradient ascent as a function of the



**Figure 13:** (a) Execution times for NNGA<sup>+</sup> with respect to the number of buckets ( $M_1$ ). (b) Execution times for NNGA<sup>+</sup> with respect to the number of neighbor layers ( $p$ ).

number of slaves and data points we can be confident in the scalability of our approach.



**Figure 14:** (a) Execution times for  $\epsilon$ -proximity with respect to the number of slaves. (b) Execution times for cluster labeling with respect to the dataset size  $n$ .

## 6 Conclusion

In this paper, we have introduced multiple improvements to the standard nearest neighbors gradient ascent used in Mean Shift algorithm. The first series of improvements are based on new usages of Locality Sensitivity Hashing for approximate nearest neighbors during the nearest neighbors gradient ascent (NNGA<sup>+</sup>), and also during cluster labeling ( $\epsilon$ -proximity).

The second one is an efficient and scalable implementation of our ideas on a distributed computing ecosystem based on Spark/Scala. We show that using our NNGA<sup>+</sup> algorithm, as a pre-processing step in other clustering methods, can improve quality metric evaluations. We demonstrated that these improvements greatly decrease the execution time whilst maintaining a suitable quality of clustering. These improvements open the opportunity to apply our Mean Shift model for Big Data clustering.

Future work is required to demonstrate the usefulness of NNGA<sup>+</sup> for clustering algorithms with very high dimensional datasets. Further study on the impact of the chosen dissimilarity measure on the  $\varepsilon$ -proximity clustering are also required and will be tackled in the future. Optimal choices of the most important tuning parameters for our proposed methods for distributed clustering, namely the number of nearest neighbors for the density gradient ascent, the number of buckets for the LSH, and the threshold for  $\varepsilon$ -proximity cluster labeling will be a subject of further investigations. Furthermore, we also desire to experiments others hashing techniques [26, 27] from the literature in a generic way in order to handle various use case.

## 7 Code

In order to facilitate further experiments and reproducible research, we will provide our contributions through an open source API which will contain NNGA<sup>+</sup>,  $\varepsilon$ -proximity algorithms and traditional Mean-shift,  $k$ -means with Spark/Scala on the following link <https://github.com/Clustering4Ever/Clustering4Ever>.

## Acknowledgments

Experiments presented in this paper were carried out using the Grid'5000 testbed, supported by a scientific interest group hosted by Inria and including CNRS, RENATER and several Universities as well as other organizations (see <https://www.grid5000.fr>).

## References

- [1] K. Fukunaga, L. Hostetler, Optimization of  $k$ -nearest-neighbor density estimates, *IEEE Trans. Inform. Theory* 19 (1973) 320–326.
- [2] K. Fukunaga, L. Hostetler, The estimation of the gradient of a density function, with applications in pattern recognition, *IEEE T. Inform. Theory* 21 (1975) 32–40.
- [3] Y. Cheng, Mean shift, mode seeking, and clustering, *IEEE Trans. Pattern Anal. Mach. Intell.* 17 (1995) 790–799.



- [4] P. Indyk, R. Motwani, Approximate nearest neighbors: Towards removing the curse of dimensionality., in: Proceedings of the 30th Annual ACM Symposium on Theory of Computing, 1998, pp. 604–613.
- [5] P. I. M. Datar, N. Immorlica, V. S. Mirrokni, Locality-sensitive hashing scheme based on p-stable distributions, in: Proceedings of the 20th Annual Symposium on Computational Geometry, 2004, pp. 253–262.
- [6] M. Slaney, M. Casey, Locality-sensitive hashing for finding nearest neighbors, in: IEEE Signal Proc. Mag., 2008, pp. 128–131.
- [7] J. Dean, S. Ghemawat, [Mapreduce: Simplified data processing on large clusters](#), Commun. ACM 51 (1) (2008) 107–113. doi:10.1145/1327452.1327492.  
URL <http://doi.acm.org/10.1145/1327452.1327492>
- [8] M. P. Forum, Mpi: A message-passing interface standard, Tech. rep., Knoxville, TN, USA (1994).
- [9] M. Zaharia, M. Chowdhury, T. Das, A. Dave, J. Ma, M. McCauly, M. J. Franklin, S. Shenker, I. Stoica, Resilient distributed datasets: A fault-tolerant abstraction for in-memory cluster computing, in: Proceedings of the 9th USENIX Symposium on Networked Systems Design and Implementation, NSDI 2012, San Jose, CA, USA, April 25-27, 2012, 2012, pp. 15–28.
- [10] C. M. Bishop, Pattern Recognition and Machine Learning, Springer-Verlag, New York, 2006.
- [11] C. C. Aggarwal, C. K. Reddy, Data Clustering: Algorithms and Applications, 1st Edition, Chapman & Hall/CRC, 2013.
- [12] B. Bahmani, B. Moseley, A. Vattani, R. Kumar, S. Vassilvitskii, Scalable K-means++, Proc. VLDB Endow. 5 (2012) 622–633.
- [13] Y. He, H. Tan, W. Luo, S. Feng, J. Fan, MR-DBSCAN: a scalable MapReduce-based DBSCAN algorithm for heavily skewed data, Front. Comp. Sc. 8 (2014) 83–99.
- [14] A. Vedaldi, S. Soatto, Quick shift and kernel methods for mode seeking, in: D. Forsyth, P. Torr, A. Zisserman (Eds.), Proceedings Part IV of the 10th European Conference on Computer Vision 2008, Marseille, France, 2008, pp. 705–718.
- [15] K.-L. Wu, M.-S. Yang, Mean shift-based clustering, Pattern Recogn. 40 (2007) 3035–3052.

- [16] D. Schugk, A. Kummert, C. Nunn, [Adaptation of the mean shift tracking algorithm to monochrome vision systems for pedestrian tracking based on HoG-features](#), Tech. rep. (2014).  
URL <https://doi.org/10.4271/2014-01-0170>
- [17] G. Z. Y. Cui, K. Cao, F. Zhang, An adaptive mean shift algorithm based on LSH, *Procedia Engineering* (2011) 265–269.
- [18] P. I. S. Har-Peled, R. Motwani, Approximate nearest neighbor: Towards removing the curse of dimensionality, *Theory Comput.* (2012) 321–350.
- [19] A. Ultsch, Clustering with SOM: U\*C, in: *Proceedings of the Workshop on Self-Organizing Maps*, 2005, pp. 75–82.
- [20] P. Fränti, Clustering basic benchmark, <http://cs.uef.fi/sipu/datasets> (2015).
- [21] L. Danon, A. Díaz-Guilera, J. Duch, A. Arenas, [Comparing community structure identification](#), *J. Stat. Mech.: Theory E*. 2005 (09) (2005) P09008.  
URL <http://stacks.iop.org/1742-5468/2005/i=09/a=P09008>
- [22] J. M. Santos, M. J. Embrechts, On the use of the adjusted rand index as a metric for evaluating supervised classification, in: C. Alippi, M. M. Polycarpou, C. G. Panayiotou, G. Ellinas (Eds.), *ICANN* (2), Springer, 2009, pp. 175–184.
- [23] D. Comaniciu, An algorithm for data-driven bandwidth selection, *IEEE T. Pattern Anal.* 25 (2003) 281–288.
- [24] W. K. Pratt, *Digital Image Processing: PIKS Inside*, 3rd Edition, John Wiley and Sons, New York, 2001.
- [25] D. Comaniciu, P. Meer, Mean shift: a robust approach toward feature space analysis, *IEEE T. Pattern Anal.* 24 (2002) 603–619.
- [26] J. Wang, H. T. Shen, J. Song, J. Ji, [Hashing for similarity search: A survey](#), *CoRR* abs/1408.2927. [arXiv:1408.2927](https://arxiv.org/abs/1408.2927).  
URL <http://arxiv.org/abs/1408.2927>
- [27] A. Morvan, A. Souloumiac, K. Choromanski, C. Gouy-Pailler, J. Atif, On the needs for rotations in hypercubic quantization hashing (2018). [arXiv:1802.03936](https://arxiv.org/abs/1802.03936).