

U*C: Self-organized Clustering with Emergent Feature Maps

Alfred Ultsch
Databionics Research Group,
University of Marburg, Germany
ultsch@informatik.uni-marburg.de

Abstract. A new clustering algorithm based on grid projections is proposed. This algorithm, called U*C, uses distance information together with density structures. The number of clusters is determined automatically. The validity of the clusters found can be judged by the U*-Matrix visualization on top of the grid. A U*-Matrix gives a combined visualization of distance and density structures of a high dimensional data set. For a set of critical clustering problems it is demonstrated that U*C clustering is superior to standard clustering algorithms such as K-means and hierarchical clusterings.

1 Introduction

Projections of high dimensional data onto two dimensions have the advantage that the structures of the data can be visually inspected. Some projections, in particular the neighborhood preserving projections, like SOM, can be used to visualize cluster structures of the high dimensional data space. We assume here a projection from data space onto a finite grid of points. ESOM (Ultsch 1999) are such a projection. Other projections like PCA, ICA or MDS can also be used if their output is discretized to a grid. ESOM have the property to disentangle cluster structures that are linear not separable. The ChainLink example (see Fig. 1) was among the first to illustrate this Ultsch/Vetter (1994). Distance relationships in a high dimensional data space can be visualized on top of a grid in form of a U-Matrix Ultsch (1990). The recently introduced P-Matrix allows the visualization of density structures of the high dimensional data space, Ultsch (2003). In this paper we present a combined visualization of distance and density in form of the U*-Matrix (Chapter 2). Then we define a novel clustering algorithm, called U*C, which uses distance and density information (Chapter 3). U*C is tested on some crucial clustering problems (Chapter 4). Its performance is compared to K-means and popular hierarchical clustering algorithms (Chapter 5). The validity and precision of the resulting clusters can be judged using the U*-Matrix. The performance of U*C is discussed in chapter 6.





Figure 1: a) ChainLink /Ward



b) U-Matrix showing disentangled clusters

2 Visualizations of Distance and Density

Let $m: D \rightarrow G$ be a mapping from a high dimensional data space $D \subset \mathbb{R}^n$ onto a finite set of positions $M = \{n_1, \dots, n_k\} \subset \mathbb{R}^2$ arranged on a grid. Each position has its two dimensional coordinates and a weight vector $W = \{w_1, \dots, w_k\}$ which is the image of a Voronoi region in D : the data set $E = \{x_1, \dots, x_d\}$ with $x_i \in D$ is mapped to a position in M such that a data point x_i is mapped to its bestmatch $bm(x_i) = n_b \in M$ with $d(x, w_b) \leq d(x, w_j) \forall w_j \in W$, where d is the distance on the data set. The set of immediate neighbors of a position n_i on the grid is denoted by $N(i)$. We call m a grid projection if m has minimal, or at least tolerable, topological errors. Compare Ultsch/Herrmann (2005) for grid dimensions and surface structures to minimize projection errors. The U-height for each position n_i is the average distance of n_i 's weight vectors to the weight vectors of its immediate neighbors $N(i)$. The U-height $uh(i)$ is calculated as follows: $uh(i) = \frac{1}{n} \sum_j d(w_i, w_j)$, $j \in N(i)$, $n = |N(i)|$.

A display of all U-heights on top of the grid is called a U-Matrix: Ultsch (1990). A single U-height shows the local distance structure. The local average distance at w_i is shown at position n_i . The overall structure of densities emerges, if a global view of a U-Matrix is regarded. Figure 1b shows an example of a U-Matrix on an ESOM with 50x82 positions. The P-height $ph(i)$ for a position n_i is a measure of the density of data points in the vicinity of w_i : $ph(i) = |\{x \in E \mid d(x, w_j) < r > 0, r \in \mathbb{R}\}|$. A display of all P-heights on top of the grid G is called a P-Matrix, Ultsch (2003). Figure 3b, 4b and 5b show examples of P-Matrices. The P-height is the number of data points within a hypersphere of radius r . The radius r should be chosen such that $ph(i)$ approximates the

probability density function of the data points. The usage of the ParetoRadius, as described in Ultsch (2003), is a choice for r . Median filtering within the $N(i)$ window can be applied on the P-Matrix. This reduces local fluctuations (noise) in the density estimation without disturbing the overall picture. Such a filtering preserves, however, the density gradients important for clustering.

For the identification of clusters in data sets it is sometimes not enough to consider distances between the data points. Consider, for example, the TwoDiamonds data set depicted in Figure 2. The data consists of two clusters of points on a plane. Inside each “diamond” the values for each data point were drawn independently from uniform distributions. At the central region, marked with an arrow in Figure 2, the distances between the data points are very small. For distance based cluster algorithms it is hard to detect correct boundaries for the clusters. Distance oriented clustering methods such as single linkage, complete linkage, Ward etc. produce classification errors. The picture changes, however, when the data’s density is regarded (see Figures 2, 3 and 5). The density at the touching point of the two diamonds is only half as big as the densities in the center regions of the clusters.




Figure 2: TwoDiamonds data set




Figure 3: a) U-Matrix of TwoDiamonds




Figure 3: b) P-Matrix of TwoDiamonds

As the TwoDiamonds data set shows, a combination of distance relationships and density relationships is necessary to give an appropriate clustering. The combination of a U-Matrix and a P-Matrix is called U*-Matrix. The U*-height $u^*h(i)$ for a position n_i is the U-height multiplied with the probability that the local density, as measured by $ph(i)$, is low. As an estimate of this probability the empirical density function can be used:

$$p_{low}(i) = Pr(\text{data density is low for position } n_i) \cong \frac{|\{p \in P\text{-matrix} \mid p > ph(i)\}|}{|\{p \in P\text{-matrix}\}|} (i).$$

The U*-height is then calculated as $u^*h(i) = uh(i) \cdot p_{low}(i)$. If the local data density is low: $u^*h(i) = uh(i)$. This happens at the presumed border of clusters. If the data density is high, then $u^*h(i) = 0$. This is in the central regions of clusters. For positions with median density holds: $u^*h(i) = uh(i) \cdot 0.5$. The U*-Matrix exhibits therefore the local data distances as heights, when the data density is low (cluster border). If the data density is high, the distances are scaled down to zero (cluster center). An alternative to formula (i) is, to adjust the multiplication factor such that $u^*h(i) = uh(i)$ for median P-heights and $u^*h(i) = 0$ for the top 20 percent of P-heights Ultsch (2003).

The cluster structure of the data can be seen more clearly on the U*-Matrix than on the U-Matrix. Figure 4 compares, for example, the U-Matrix taken from Kaski et al (1999) to a U*-Matrix of the same data set. Since density and distance play different roles in the definition of clusters, we think that the three different matrices, U-, P- and U*-Matrix together give an appropriate impression of the cluster structure of any high dimensional data (see Figure 4). For the U-, P- and U*-Matrix the grid (also called the map) is the floor space layout for a landscape like visualization of distance- and density structures of the high dimensional data space. Structures emerge on top of the grid by regarding an overall view of many positions. Single positions are only tiny parts of these structures (Ultsch (1999)). The U-, U*- and P- Matrix visualizations of the data are called emergent grid projections (EGP).





Figure 4: a) U-Matrix/ Kaski et al (1999),



b) U*-Matrix, same data.

3 U*C Clustering Algorithm

A topological correct mapping m projects a cluster onto a coherent area on the grid (cluster area). Points within the cluster are mapped to the inside of the cluster area. Data points at the border (surface) of the cluster are projected to the border of the cluster area. Consider a data point x at the surface of a cluster C , with $n_i = bm(x)$. The weight vectors of its neighbors $N(i)$ are either within the cluster, in a different cluster or interpolate between clusters. If we assume that the inter cluster distances are locally larger than the local inner cluster distances, then the U-heights in $N(i)$ will be large in such directions which point away from the cluster C . This means, a gradient descent on the U-Matrix will lead away from cluster borders. A movement from one position n_i to another position n_j with the result that w_j is more within a cluster C than w_i is called immersive. For data points well within C , a gradient descent on a U-Matrix will, however, not necessarily be immersive. The P-heights follow the density structure of a cluster. Under the assumption that the core parts of a cluster are those regions with largest density, a gradient ascent on the P-Matrix is immersive. Clusters may also be defined by density alone instead of distance. See, for example, the EngyTime data set shown in Figure 5a and its density structure as shown by the P-Matrix in Figure 5b. This data set represents situations where the data generation can be described appropriately by Hidden Markov Models.

At the borders of a cluster the measurement of density is, however, critical. At cluster borders the local density of the points should decrease substantially. In most cases the cluster borders are defined either by low point densities (see Figure 5) or by “empty space” between clusters (= large inter cluster distances). For empirical estimates of the point density a gradient ascent on a P-Matrix may therefore not be immersive for points at cluster

borders. A movement on a grid which follows first a gradient descent on a U-Matrix and then a gradient ascent on a P-Matrix is called immersion. Let I denote the end points of immersion starting from every position on a grid. If the density within a cluster is constant, immersion will not converge to a single point for a cluster for all starting points within a cluster. The U*-Matrix is then used to determine which points in I belong to the same cluster. The watersheds of the U*-Matrix are calculated using the algorithm described in Luc/Soille (1991). Points that are separated by a watershed are assigned to different clusters, points within the same basin to a single cluster. The following pseudocode summarizes the U*C clustering algorithm described above.




Figure 5: a) EngyTime/SingleLinkage clustering




Figure 5: b) P-Matrix of EngyTime

U*C clustering Algorithm: given U-Matrix, P-Matrix, U*-Matrix, $I = \{ \}$;

Immersion:

For all positions n of the grid:

- 1) from position n follow a gradient descent on the U-Matrix until a minimum is reached in position u
- 2) from position u follow a gradient ascent on the P-Matrix until a maximum is reached in position p .
- 3) $I = I \cup \{p\}$; Immersion(n) = p .

Cluster assignment:

- 1) calculate the watersheds for the U*-Matrix (e.g. using Luc/Soille (1991)).
- 2) partition I using these watersheds into clusters C_1, \dots, C_c
- 3) assign a data point x to a cluster C_j if Immersion(bm(x)) $\in C_j$.

4 Fundamental Clustering Problems Suite

The efficiency of the U*C clustering algorithm is tested using a set of ten clustering problems called Hepta, Lsun, Tetra, Chainlink, Atom, EngyTime, Target, TwoDiamonds, WingNut and GolfBall. Any reasonable clustering algorithm should be able to solve these problems correctly Jain/Dubes (1998). As can be seen below, however, standard algorithms like K-means, and hierarchical clustering algorithms, like single linkage and Ward have difficulties on several data sets. The suite of data sets is called Fundamental Clustering Problem Suite (FCPS). The suite can be downloaded from the website of the author: <http://www.informatik.uni-marburg.de/~databionics>.

FCPS poses some hard clustering problems. Chainlink and Atom are not separable by hyperplanes. The GolfBall data set consists of points that are equidistant on the surface of a sphere. This data set is used to address the problem to impose cluster structures when no such structure is present. The problem of outliers is addressed by the Target data set shown in Figure 6.




Figure 6: SL, Ward and K-means clustering algorithms on the Target data set

The following is a short description of the data set and the problem it poses to cluster algorithms. Pictures of the data sets are shown in the Figures of this paper.

Name	n	d	k	Main Problem
Hepta	212	3	7	different densities
Lsun	400	2	3	different variances
Tetra	400	3	4	inner vs inter dist.
Chainlink	1000	3	2	not linear separable
Atom	800	3	2	lin. not sep., diff. dens.
EngyTime	4096	2	2	density defined clusters
Target	770	2	6	outliers
TwoDiamonds	800	2	2	touching clusters
WingNut	1070	2	2	densities at borders
GolfBall	4002	3	1	no cluster at all

5 Results

The results of U*C Clustering are compared to K-means as most popular partitioning cluster algorithm. The hierarchical cluster algorithms SingleLinkage and Ward were applied to the FCPS data sets. All algorithms were used as implemented in MATLAB™. Since K-means converges to a local minimum of a cost function, the best of 100 repetitions was kept. The correct number of clusters was given as parameter to all standard algorithms. Shown are the overall accuracies. Performances lower than 80% are emphasized. There is no data set on which U*C performs worse than any of the other clustering algorithms.

Data Set	Single	Ward	K-means	U*C
Hepta	100 %	100 %	100 %	100 %
Lsun	100 %	50 %	50 %	100 %
Tetra	0.01 %	90 %	100 %	100 %
Chainlink	100 %	50 %	50 %	100 %
Atom	100 %	50 %	50 %	100 %
EngyTime	0 %	90 %	90 %	90 %
Target	100 %	25 %	25 %	100 %
TwoDiamonds	0 %	100 %	100 %	100 %
WingNut	0 %	80 %	80 %	100 %
GolfBall	100 %	-	-	100 %

Table 1: Accuracy of the clustering algorithms on FCPS

6 Discussion

SingleLinkage (SL) clustering imposes a chain of data points as cluster model on the data. This algorithm is usually misled, if the local inner cluster distances are in the same range as the inter cluster distances. This can be seen in the performance on the Tetra, EngyTime, Two Diamonds and WingNut data set. SL is, however, able to separate clusters that are not separable by hyperplanes, e.g. ChainLink and Atom. The fundamental cluster model of Ward is a hyperellipsoid. For data sets which do not fit this model (5 of the 10 sets in FCPS), Ward produces an erroneous clustering. Most clustering algorithms require the knowledge of the number of clusters. U*C determines the number of clusters automatically. It is, however, crucial to assess the validity of a clustering. Dendrograms as used in hierarchical clustering

algorithms seem to address this, are, however, sometimes misleading as the following example shows. Figure 7a shows a dendrogram of the GolfBall data using Ward hierarchical clustering. Such a dendrogram would suggest 3 or 6 cluster.

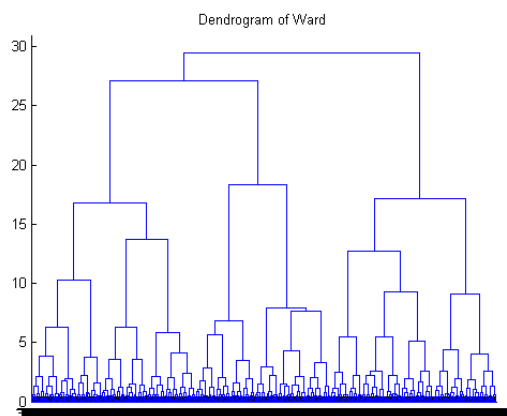


Figure 7: a) GolfBall Dendrogram of Ward




Figure 7: b) Golf Ball / K-means

K-means requires that the clusters are compact and have about the same variance. The boundary separating two clusters in K-means is always a hyperplane. This results in the bad performance of K-means on 5 of the 10 data sets in the FCPS. K-means enforces any number of clusters on the data. In Figure 7b a K-means clustering of GolfBall with $k = 6$ is shown. For high dimensional data this result may, however, be taken serious, in particular in the light of the dendrogram of Figure 7a. The U*-Matrix serves as a feedback of the cluster structure in U*C. Compare Figure 4b, for example, to see, whether there are cluster structures in a very high dimensional data set.

7 Conclusion

A new clustering algorithm based on grid projections is proposed. This algorithm uses distance structures (U-Matrix) as well as density structures (P-Matrix) of the data set. No particular geometrical cluster model is imposed on the data by U*C. Other clustering algorithms impose such a model and are performing poor, if the data set is of a different structure. The

number of clusters is determined automatically in U*C. The correctness and validity of the clusters found can be assessed directly using the U*-Matrix visualization. The U*-Matrix shows a combined picture of distance and density structures of a high dimensional data set. U*C performs superior to standard clustering algorithms such as K-means and the most popular hierarchical algorithms Jain/Dubes (1998). This is demonstrated on a group of data sets which represent fundamental clustering problems, like different variances, outliers and other structural difficulties. U*C and other tools for ESOM, see Ultsch/Mörchen (2005), can be downloaded from our web site.

References

- A. ULTSCH, A., MÖRCHEN, F. (2005): ESOM-Maps: tools for clustering, visualization, and classification with Emergent SOM, Dept. of Computer Science University of Marburg, Research Report. 46.
- JAIN, A.K., DUBES, R.C. (1998): Algorithms for Clustering Data, New York, Wiley.
- KASKI ET AL (1999): Analysis and Visualisation of Gene Expression Data using Self Organizing Maps, Proc NSIP.
- LUC, V, SOILLE, P. (1991): Watersheds in Digital Spaces: An Efficient Algorithm Based on Immersion Simulations, IEEE Transactions of Pattern Analysis and Machine Intelligence, Vol. 13(6), 583-598.
- ULTSCH, A., VETTER, C. (1994): Selforganizing Feature Maps vs. statistical clustering, Dept. of Computer Science University of Marburg, Research Report 9.
- ULTSCH, A. (2003): Maps for the Visualization of high-dimensional Data Spaces, In Proc. WSOM, Kyushu, Japan, 225-230.
- ULTSCH, A., (1999): Data Mining and Knowledge Discovery with Emergent Self-Organizing Feature Maps for Multivariate Time Series, E. Oja and S. Kaski (eds), Kohonen Maps, 33-46.
- ULTSCH, A., (2003): U*-Matrix: A Tool to visualize Clusters in high dimensional Data, Dept. of Computer Science University of Marburg, Research Report 36.
- ULTSCH, A., HERRMANN, L. (2005), The architecture of Emergent Self-Organizing Maps to reduce projection errors, Proc ESANN, Brugges 2005.
- ULTSCH, A., SIEMON, H.P. (1990): Kohonen's Self Organizing Feature Maps for Exploratory Data Analysis, In Proc. Intern. Neural Networks, Kluwer Academic Press, Paris, 305-308.