

A Novel Construction of Connectivity Graphs for Clustering and Visualization

Wesam Barbakh
University of the West of Scotland
School of Computing
Paisley, PA1 2BE, Scotland
UK
wesam.barbakh@uws.ac.uk

Colin Fyfe
University of the West of Scotland
School of Computing
Paisley, PA1 2BE, Scotland
UK
colin.fyfe@uws.ac.uk

Abstract: We [5, 6] have recently investigated several families of clustering algorithms. In this paper, we show how a novel similarity function can be integrated into one of our algorithms as a method of performing clustering and show that the resulting method is superior to existing methods in that it can be shown to reliably find a globally optimal clustering rather than local optima which other methods often find. We discuss some of the current difficulties with using connectivity graphs for solving clustering problems, and then we introduce a new algorithm to build the connectivity graphs. We compare this new algorithm with some famous algorithms used to build connectivity graphs. The new algorithm is shown to be superior to those in the current literature. We also extend the method to perform topology preserving mappings and show the results of such mappings on artificial and real data.

Key-Words: Clustering, Similarity function, Connectivity graph, Visualisation, K-means, Topographic mapping.

1 Introduction

One of the major tasks falling to a data analyst is to find groups of data all of which share some underlying feature which is not shared by other samples in the dataset. This is known as clustering the data if it is done in an unsupervised manner i.e. if we do not give the method any class information on which to work. The K-means algorithm [16, 20, 23] is one of the most frequently used investigatory algorithms in data analysis. The algorithm attempts to locate K prototypes or means throughout a data set in such a way that the K prototypes in some way best represents the data. It is an iterative algorithm in which K means are spread throughout the data and the data samples are allocated to the mean which is closest (often in Euclidean norm) to the sample. Then the K means are repositioned as the average of data points allocated to each mean. This continues until stable convergence is reached. The K-means algorithm is one of the first which a data analyst will use to investigate a new data set because it is algorithmically simple, relatively robust and gives ‘good enough’ answers over a wide variety of data sets: it will often not be the single best algorithm on any individual data set but it may be close to the optimal over a wide range of data sets. However the algorithm is known to suffer from the defect that the means or prototypes found depend on the initial values given to them at the start

of the simulation: a typical program will converge to a local optimum. There are a number of heuristics in the literature which attempt to address this issue but, at heart, the fault lies in the performance function on which K-means is based.

[19] proposed a global K-means algorithm, an incremental approach to clustering that adds one cluster prototype at a time through a deterministic global search consisting of N (the data size) executions of the K-means; this algorithm can obtain equivalent or better results than the standard K-means, but it suffers from high computation cost and at the same time gives no guarantee to find the optimum.

Arthur and Vassilvitskii [2] improved the K-means algorithm by substituting the random allocation of the prototypes with a seeding technique. They give experimental results that show the advantage of this algorithm in time and accuracy.

A variation on K-means is the so-called soft K-means [22] in which prototypes are allocated according to

$$\mathbf{m}_k = \frac{\sum_n r_{kn} \mathbf{x}_n}{\sum_{j,n} r_{jn}} \quad (1)$$

$$\text{where e.g. } r_{kn} = \frac{\exp(-\beta d(\mathbf{x}_n, \mathbf{m}_k))}{\sum_j \exp(-\beta d(\mathbf{x}_n, \mathbf{m}_j))} \quad (2)$$

and $d(a, b)$ is the Euclidean distance between a and b .




Figure 1: The vagaries of simulations: top, the soft K-means failed to identify the four clusters but on the bottom is successful with the same data set.

β is a “stiffness” parameter. Note that the standard K-means algorithm is a special case of the soft K-means algorithm in which the responsibilities, $r_{kn} = 1$ when \mathbf{m}_k is the closest prototype to \mathbf{x}_n and 0 otherwise. However the soft K-means does increase the non-localness of the interaction since the responsibilities are typically never exactly equal to 0 for any data point-prototype combination.

However there are still problems with soft K-means. We find that with soft K-means it is important to choose a good value for β ; if we choose a poor value we may have poor results in finding the clusters. However, even with a good value, we often still find that soft K-means has the problem of sensitivity to the prototypes’ initialization. As shown in Fig. 1, while soft K-means succeeds in identifying the clusters (bottom diagram), sometimes it failed for the same data sample when we used a different initialization of the prototypes (top diagram).

In [3, 5, 26] we derive a family of new clustering algorithms that solve the problem of sensitivity to initial conditions in the K-means algorithm. In this paper, we show how to extend one of our algorithms, the inverse weighted clustering IWC [7, 8], to do clustering based on similarity functions. We take as an

example the exponential function to measure the similarity between data points and prototypes; we can use other similarity functions. We compare the new algorithm with the soft K-means algorithm as it is also based on the exponential function. We show in this paper that the new algorithm gives better results than both K-means and soft K-means. Also, this new algorithm solves the problem of excessive computation in [19].

2 Clustering with similarity functions

The inverse weighted clustering algorithm (IWC) has the following rule:

$$\mathbf{m}_k = \frac{\sum_{i=1}^N b_{ik} \mathbf{x}_i}{\sum_{i=1}^N b_{ik}} \quad (3)$$

where

$$b_{ik} = \frac{\|\mathbf{x}_i - \mathbf{m}_{k*}\|^{P+2}}{\|\mathbf{x}_i - \mathbf{m}_k\|^{P+2}} \quad (4)$$

(3) introduces the clustering algorithm (IWC). We have given extensive simulations in [7] showing that this algorithm is insensitive to the prototypes’ initialization. The question now, is how can we use this algorithm for clustering if we have similarity functions that measure the similarity between objects (or data points)? How can we work with the similarity function and at the same time get the benefit of (3)?

2.1 Exponential function as similarity function

Let the exponential function be used, as an example, to measure the similarities between points as similarity(\mathbf{x}, \mathbf{m}) = $\exp(-\|\mathbf{x} - \mathbf{m}\|)$.

To use this similarity function for clustering while taking the benefits of (3) we need to go through the following steps:

1. Measure the similarities between all data points and prototypes.
2. For alternative similarity function, if the similarity measurements are outwith [0,1], normalize, so that 1 corresponds to the highest similarity and 0 corresponds to no similarity.
3. Map the similarity measurements to distances: Distance (\mathbf{x}, \mathbf{m}) = $1 - \exp(-\|\mathbf{x} - \mathbf{m}\|)$.
4. Compute (3) for all prototypes using the new distance measurements generated from the similar-




Figure 2: Artificial data set: 7 clusters of red '*'s, and 7 prototypes of blue 'o's.

ity function.

$$b_{ik} = \frac{\text{distance}(\mathbf{x}_i - \mathbf{m}_{k*})^{P+2}}{\text{distance}(\mathbf{x}_i - \mathbf{m}_k)^{P+2}} \quad (5)$$

We will call this Inverse Weighted Clustering with Similarity Function algorithm (IWCwSF)[4, 9].

2.2 Simulations

Fig. 2 shows an artificial data set consisting of 7 clusters of red '*'s, and all the prototypes are initialized to lie within one cluster and are shown as blue 'o's. Fig. 3 shows the results after applying soft K-means and inverse weighted clustering with similarity function (IWCwSF) to this artificial data set shown in Fig. 2. In Fig. 3, top, we see the first cluster, bottom left, is divided into three sub clusters. Also, the upper two clusters, left, are grouped together as one cluster. This poor result from the soft K-means appeared due to its convergence to a local optimum. As shown in Figure 3, while the soft K-means failed to identify the clusters, top, the IWCwSF algorithm identified all of them successfully, bottom.

To show how the new algorithm behaves with dead prototypes, we have in Fig. 4 another artificial data set consisting of 40 data points. Each data point represents a cluster, so we have 40 clusters, and 40 prototypes are initialized randomly and very far from data to represent some dead prototypes. Fig. 5 shows the result after applying soft K-means and IWCwSF to the artificial data set shown in Fig. 4. Again, the new algorithm IWCwSF identified all the clusters successfully, bottom, while soft K-means algorithm didn't, top.

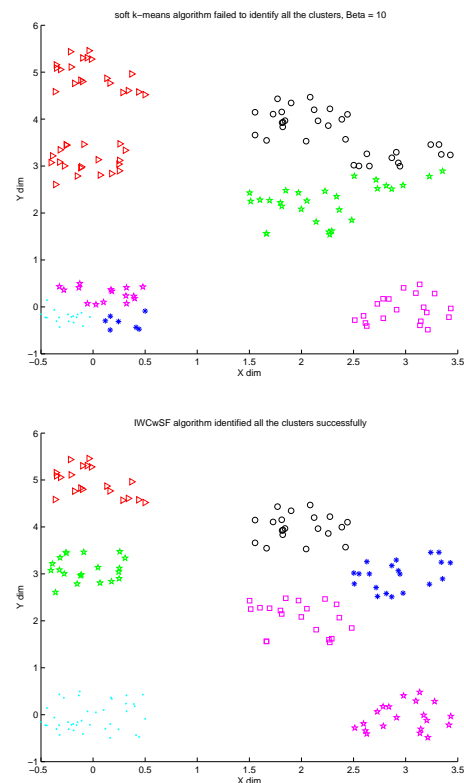


Figure 3: Top: result after applying soft K-means algorithm. Bottom: result after applying inverse weighted clustering with similarity function (IWCwSF).




Figure 4: Artificial data set: 40 clusters of red '*'s, and 40 prototypes of blue 'o's.

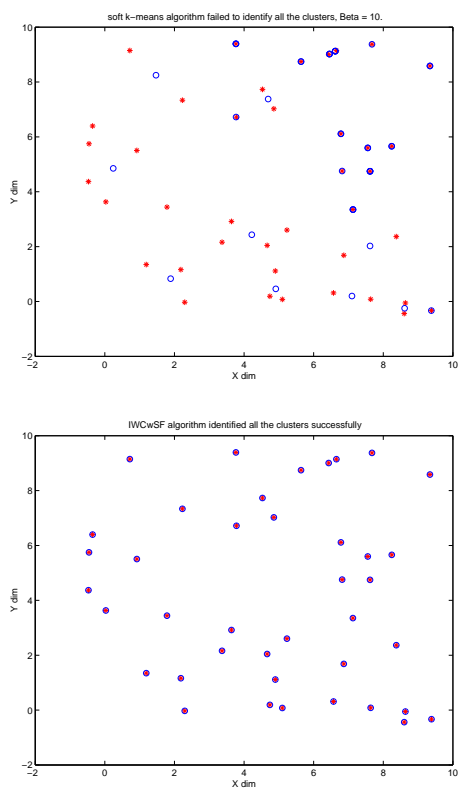


Figure 5: Top: result after applying soft K-means algorithm to the data of Fig. 4. Bottom: result after applying inverse weighted clustering with similarity function (IWCwSF).

3 Different connectivity graphs

To construct a connectivity graph we transform a given set x_1, \dots, x_n of data points with pairwise similarities s_{ij} or distances d_{ij} into a graph. There are several popular methods to construct similarity graphs [21]. The goal of constructing connectivity graphs is to model the local neighborhood relationships between data points. In this section we review two popular methods to construct a connectivity graph, and then we introduce a new algorithm that solves some problems that can not be solved by the others.

3.1 The ε -neighborhood graph

In this algorithm we connect all points whose pairwise distances are smaller than ε . Choosing a useful value of ε is not a trivial task [21]. Sometimes it is not possible to find a suitable value of ε if we have data on different scales or densities.

3.2 k-nearest neighbor graphs

In this algorithm we connect each point to all points among its k-nearest neighbors. This leads to a directed graph as the neighborhood relationship is not necessarily symmetric. To make this graph undirected we have two choices. The first one is to simply ignore the direction of the edges, so that points x and y are connected with undirected edge if x is among the k-nearest neighbors of y , or if y is among the k-nearest neighbors of x . The resulting graph is usually called k-nearest neighbor graph. The second way is to connect points x and y if both of them are among the k-nearest neighbors of the other. The resulting graph is called mutual k-nearest neighbor graph [21].

3.3 New connectivity graph

We show here a new connectivity graph that can be used for clustering and spectral clustering [25, 28]. In this algorithm, we connect all points whose pairwise weights are higher than a useful threshold. This threshold is a value ranging between 0 and 1; if its value equals 0, all the points will be connected together. If its value equals 1, only the points which are close to each other will be connected. The weight between vertex \mathbf{v}_i and \mathbf{v}_j is calculated by

$$weight(\mathbf{v}_i, \mathbf{v}_j) = \frac{\min(X, Y)}{\|\mathbf{v}_i - \mathbf{v}_j\|} \quad (6)$$

where X is the distance between \mathbf{v}_i and the closest vertex to it.

Y is the distance between \mathbf{v}_j and the closest vertex to it.




Figure 6: Artificial data set is shown as 2 clusters of red '*'s.

(6) has the same motivation as our clustering algorithms in [5, 7, 8], see (4) which is used to measure the similarity between a data point and a prototype. In (6), we use the same idea to measure the similarity between two data points.

This new connectivity graph algorithm solves some problems that can not be solved by the previous algorithms as we will see in the next example. However this new algorithm is not designed to replace the previous algorithms but to add new behavior to them and to give better results for different applications.

Example 1:

Suppose that we have an artificial data set like that shown in Figure 6. x_1 and x_2 represent the first diffuse cluster, and x_3 and x_4 represent the second tighter cluster. Figure 7 shows the results after applying ϵ -neighborhood graph, with different values of ϵ ($\epsilon=0.5, 1.1, 1.5$ and 2.1 respectively), to the artificial data set shown in Figure 6. From Figure 7 we can see it is not possible to find a useful parameter ϵ that defines the two clusters successfully. Figure 8 shows the results after applying k-nearest neighbor graphs, with $k=1, 2$ and 3 , to the artificial data set shown in Figure 6. From Figure 8 we can see there is no useful parameter k that defines the two clusters successfully. Figure 9 shows the results after applying mutual k-nearest neighbor graphs, with $k=1$ and 2 , to the artificial data set shown in Figure 6. Figure 10 shows the result after applying the new connectivity graph algorithm to the artificial data set shown in Figure 6. As shown in Figure 10 the new algorithm succeeds in identifying the two clusters.

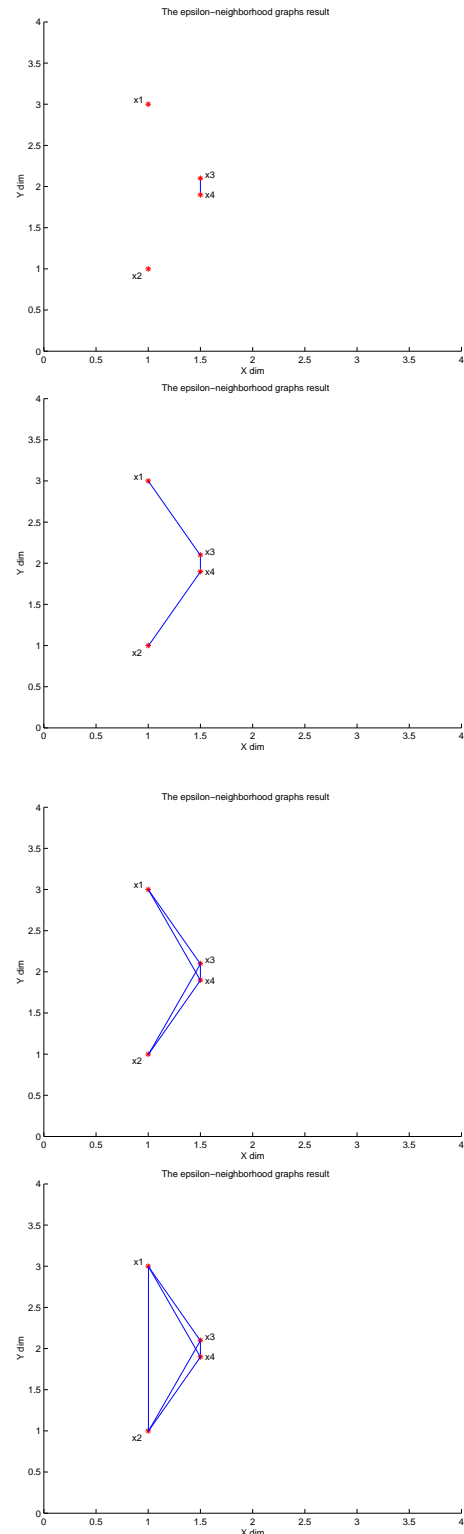


Figure 7: The ϵ -neighborhood graph results with different values of ϵ ($\epsilon=0.5, 1.1, 1.5$ and 2.1 respectively).




Figure 8: The k-nearest neighbor graphs results with k=1, 2 and 3.




Figure 9: The mutual k-nearest neighbor graphs results with k=1 and 2.

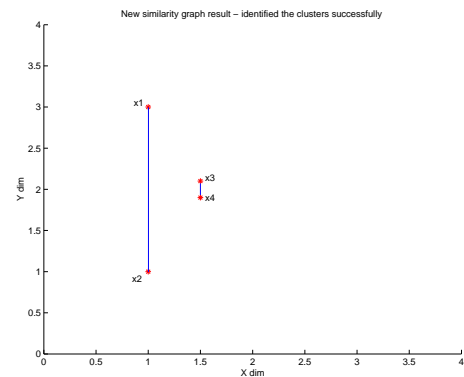


Figure 10: The new connectivity graph algorithm succeeds in identifying the clusters.




Figure 11: Artificial data set consisting of 3 circles of red '*'s. The outer and inner circles have roughly similar density.




Figure 12: The new connectivity graph algorithm results. Top: the new algorithm identifies each circle as a cluster. Bottom: the new algorithm identifies outer and inner circles as one cluster.

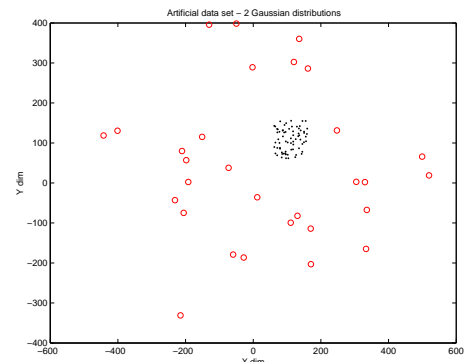


Figure 13: Artificial data set taken from two different Gaussian distributions, the red 'o's represents the first distribution and the black dots represents the second distributions.

3.4 Simulations

Figure 12 shows the results after applying the new connectivity graph algorithm, with two different threshold values, to the artificial data set shown in Figure 11. In Figure 12, top, the new algorithm identifies each circle as a cluster. We can get another result by choosing another threshold value as shown in Figure 12, bottom. In Figure 12, bottom, the outer and inner circles have roughly similar density and are identified as one cluster while the middle circle, which has higher density, is identified as another cluster. This result can not be achieved by either ϵ -neighborhood graph or k-nearest neighbor graphs.

In Figure 13, we have an artificial data set taken from two different Gaussian distributions. Figure 14 shows the results after applying the new connectivity graph algorithm to the artificial data set shown in Figure 13. From Figure 14, top, we can see that all the data points from first distribution are connected together without any connection to the points from the second distribution.

4 A Topology Preserving Mapping

In this section we show how it is possible to extend the IWCSF algorithm to provide a new algorithm for visualization and topology-preserving mappings.

4.1 Inverse Weighted Clustering with Similarity Function Topology Preserving Mapping(IWCSFToM)

A topographic mapping (or topology preserving mapping) is a transformation which captures some structure in the data so that points which are mapped close




Figure 14: The new connectivity graph algorithm results. Top: the data points in the first distribution are connected together. Bottom: final result, clusters identified successfully from the points that are connected together.

to one another share some common feature while points which are mapped far from one another do not share this feature. The Self-organizing Map (SOM) was introduced as a data quantisation method but has found at least as much use as a visualisation tool.

Topology-preserving mappings such as the Self-organizing Map (SOM) [17, 18] and the Generative Topographic Mapping (GTM) [10, 11, 12, 29, 30] have been very popular for data visualization: we project the data onto the map which is usually two dimensional and look for structure in the projected map by eye. We have recently investigated a family of topology preserving mappings [15] which are based on the same underlying structure as the GTM.

The basis of our model is K latent points, t_1, t_2, \dots, t_K , which are going to generate the K prototypes, \mathbf{m}_k . To allow local and non-linear modeling, we map those latent points through a set of M basis functions, $f_1(), f_2(), \dots, f_M()$. This gives us a matrix Φ where $\phi_{kj} = f_j(t_k)$. Thus each row of Φ is the response of the basis functions to one latent point, or alternatively we may state that each column of Φ is the response of one of the basis functions to the set of latent points. One of the functions, $f_j()$, acts as a bias term and is set to one for every input. Typically the others are gaussians centered in the latent space. The output of these functions are then mapped by a set of weights, W , into data space. W is $M \times D$, where D is the dimensionality of the data space, and is the sole parameter which we change during training. We will use \mathbf{w}_i to represent the i^{th} column of W and Φ_j to represent the row vector of the mapping of the j^{th} latent point. Thus each basis point is mapped to a point in data space, $\mathbf{m}_j = (\Phi_j W)^T$.

We may update W either in batch mode or with online learning: with the Topographic Product of Experts [15], we used a weighted mean squared error; with the Harmonic Topographic Mapping [15, 24, 27], we used Harmonic K-means [31, 32]. We now apply the IWCwSF algorithm to the same underlying structure to create a new topology preserving algorithm.

Each data point is visualized as residing at the prototype on the map which would win the competition for that data point. However we can do rather better by defining the responsibility that the j^{th} prototype has for the i^{th} data point as

$$r_{ji} = \frac{\exp(-\gamma \|\mathbf{x}_i - \mathbf{w}_j\|^2)}{\sum_k \exp(-\gamma \|\mathbf{x}_i - \mathbf{w}_k\|^2)} \quad (7)$$

We then project points taking into account these responsibilities: let y_{ij} be the projection of the i^{th} data point onto the j^{th} dimension of the latent space; then

$$y_{ij} = \sum_k t_{kj} r_{ki} \quad (8)$$




Figure 15: The resulting prototypes' positions after applying IWCSFToM. The prototypes are shown as blue 'o's.

where t_{kj} is the j^{th} coordinate of the k^{th} latent point. When we use these algorithms for visualisation purposes, it is these y-values (which are typically two dimensional coordinates) which we use. Note that this method represents each data point \mathbf{x}_i by a value y_i where y_i is a weighted sum of the coordinates of the original latent points. An alternative (which is typically used the SOM) is to find the latent point with greatest responsibility for the data point and allocate its y_i value at this latent point.

4.2 Simulations

o Artificial data set

We create a simulation with 10 latent points deemed to be equally spaced in a one dimensional latent space, passed through 5 Gaussian basis functions and then mapped to the data space by the linear mapping W which is the only parameter we adjust. We generated 500 two dimensional data points, (x_1, x_2) , from the function $x_2 = x_1 + 1.25 \sin(x_1) + \mu$ where μ is noise from a uniform distribution in $[0,1]$. Final result from the IWCSFToM is shown in Fig. 15 in which the projections of consecutive latent points are joined. We see that nearby latent points take responsibility for nearby data points.

o Real data sets

1. The iris¹ data set is a data set with 150 random samples of flowers from the iris species setosa, versicolor, and virginica collected by Anderson in 1935 [1]. There are 50 observations from each species for sepal length, sepal width, petal length and petal width in cm. This data set was used by Fisher(1936) [13] in his

initiation of linear discriminant-function technique.

2. In the algae² data set we have 72 samples classified into 9 types. Each sample is recorded as an 18 dimensional vector representing the magnitudes of various pigments.
3. The Bank³ data set appeared in [14]. This data set has 200 observations. It contains measurements on 100 forged and 100 genuine bank notes. Each data record contains the six measurements:

LENGTH: length of bill.

LEFT: width of bill, measured on the left.

RIGHT: width of bill, measured on the right.

BOTTOM: width of the margin at the bottom.

TOP: width of the margin at the top.

DIAGONAL: length of the image diagonal.

Fig. 16 shows the results of applying IWCSFToM to the iris, algae and bank data set. The real data sets are projected onto a two dimensional grid of latent points (10 x 10). For the iris data set, top, we can see one cluster is distant and separated, while there is difficulty to separate the other two completely. For the algae data set, middle, we have 7 clusters separated, and the two clusters to the left are grouped together. For the bank data set, bottom, the two clusters are identified successfully.

5 Conclusion

We have previously investigated several new families of clustering algorithms [5, 6]. In this paper, we have shown how to get the benefits from the new algorithms but now based on alternative similarity functions. We have used artificial data with deliberately poor initialization since with real data, we often do not know what is a good and what is a poor initialization. This is exacerbated by the 'curse of dimensionality': even when we decide to use means of data as cluster prototypes, these can often lie far from the actual data which may be found in the outer shell of a high dimensional sphere.

We have illustrated a new method of identifying clusters using the concepts of connectivity graphs. However unlike the standard methods, our method takes into account, not just the distance between pairs of data points, but also the shortest distance between

²The algae data set is available at <http://mlearn.ics.uci.edu/databases/>.

³The Bank data set is available at <http://www.quantlet.com/mdstat/scripts/mva/htmlbook/mvhtmlnode129.html>.

¹The iris data set is available at <http://mlearn.ics.uci.edu/databases/>.




Figure 16: The results of applying IWCSFToM on the real data sets, iris, top, algae, middle and bank, bottom.

either data point and any other data point. In doing so, we can account for variable density of data points within a data set. Initial results on artificial data sets are very promising.

Finally we have extended the proposed algorithm for visualisation by incorporating a latent or hidden space which underlies the prototypes' positions in data space. By constraining the latent points' positions to certain values, we can ensure that we have preserved local neighbourhood relations with the prototypes' positions in data space. We have illustrated the effect of this with particular emphasis on the power of our method as a visualization tool. We have shown typical results with the projections of real data into an underlying two dimensional latent space.

We will, in future, investigate alternative similarity functions to evaluate which is best for which data and under what conditions. There may be no globally optimal similarity function but there may be underlying principles which we can use to determine the optimal similarity function for a particular type of data.

References:

- [1] E. Anderson. The irises of the gaspe peninsula. *Bulletin of the American Iris Society*, 59:2–5, 1935.
- [2] D. Arthur and S. Vassilvitskii. K-means++: The advantages of careful seeding. In *Bay Area Theory Symposium, BATS 06*, 2006. <http://www.stanford.edu/~sergeiv/papers/kMeansPP-soda.pdf>.
- [3] W. Barbakh. The family of inverse exponential k-means algorithms. *Computing and Information Systems*, 11(1):1–10, February 2007. ISSN 1352-9404.
- [4] W. Barbakh. *Local versus Global Interactions in Clustering Algorithms*. PhD thesis, School of Computing, University of the West of Scotland, UK, 2008.
- [5] W. Barbakh, M. Crowe, and C. Fyfe. A family of novel clustering algorithms. In *7th international conference on intelligent data engineering and automated learning, IDEAL2006*, pages 283–290, University of Burgos, Spain, September 2006. Springer. ISSN 0302-9743 ISBN-13 978-3-540-45485-4.
- [6] W. Barbakh and C. Fyfe. Clustering with reinforcement learning. In *International Conference on Intelligent Data Engineering and Automated Learning IDEAL'07*, pages 507–516, University of Birmingham, UK, December 2007. Springer. LNCS 4881.
- [7] W. Barbakh and C. Fyfe. Inverse weighted clustering algorithm. *Computing and Information Systems*, 11(2):10–18, May 2007. ISSN 1352-9404.
- [8] W. Barbakh and C. Fyfe. Tailoring local and global interactions in clustering algorithms. Technical Report 40, School of Computing, University of the West of Scotland, University of Paisley, UK, March 2007. ISSN 1461-6122.

- [9] W. Barbakh and C. Fyfe. Clustering and visualization with alternative similarity functions. In *The 7th WSEAS international conference on artificial intelligence, knowledge, engineering and data bases, AIKED'08*, pages 238–244, University of Cambridge, UK, February 2008. ISSN: 1790-5109 ISBN: 978-960-6766-41-1.
- [10] C. M. Bishop, M. Svensen, and C. K. I. Williams. Gtm: The generative topographic mapping. *Neural Computation*, 10(1):215–234, 1997.
- [11] C. M. Bishop, M. Svensen, and C. K. I. Williams. Magnification factors for the gtm algorithm. In *Proceeding of the IEE 5th International Conference on Artificial Neural Networks, Cambridge, UK*, pages 64–69, 1997.
- [12] C. M. Bishop, M. Svensen, and C. K. I. Williams. Developments of the generative topographic mapping. *Neurocomputing*, 21(1):203–224, 1998.
- [13] R.A. Fisher. The use of multiple measurements in taxonomic problems. *Annals of Eugenics* 7, pages 179–188, 1936.
- [14] B. Flury and H. Riedwyl. *Multivariate Statistics: A practical approach*. Cambridge University Press, 1988.
- [15] C. Fyfe. Two topographic maps for data visualization. *Data Mining and Knowledge Discovery*, 14:207–224, 2007. ISSN 1384-5810.
- [16] J. Hartigan and M. Wang. A k-means clustering algorithm. *Applied Statistics*, 28:100–108, 1979.
- [17] T. Kohonen. *Self-Organization and Associative Memory*. Springer-Verlag, 1984.
- [18] T. Kohonen. *Self-Organising Maps*. Springer, 1995.
- [19] A. Likas, N. Vlassis, and J. J. Verbeek. The global k-means clustering algorithm. *Pattern Recognition* 36, pages 451–461, 2003.
- [20] S. P. Lloyd. *Least squares quantization in pcm*. Technical note, Bell Laboratories, 1957. Published in 1982 in *IEEE Transactions on Information Theory* 28, 128-137.
- [21] U. V. Luxburg. A tutorial on spectral clustering. Technical Report TR-149, Max Planck Institute for Biological Cybernetics, March 2007.
- [22] D. J. MacKay. *Information Theory, Inference, and Learning Algorithms*. Cambridge University Press., 2003.
- [23] J. MacQueen. Some methods for classification and analysis of multivariate observations. *Proceedings of the fifth Berkeley symposium on mathematical statistics and probability*, pages 281–297, 1967.
- [24] S. McGlinchey and M. Pena. Quantization errors in the harmonic topographic mapping. In *The 9th WSEAS International Conference on applied mathematics, MATH 06*, pages 105–110, May 2006.
- [25] A. Y. Ng, M. Jordan, and Y. Weiss. On spectral clustering: Analysis and an algorithm. In *Proc. of NIPS-14*, pages 849–856, 2001. <http://ai.stanford.edu/~ang/papers/nips01-spectral.ps>.
- [26] M. Pena, W. Barbakh, and C. Fyfe. *Principal Manifolds for Data Visualization and Dimension Reduction*, chapter Topology-Preserving Mappings for Data Visualisation, pages 132–152. Springer, 2007. ISBN 978-3-540-73749-0.
- [27] M. Pena and C. Fyfe. The harmonic topographic map. Technical Report 35, School of Computing, University of Paisley, 2005.
- [28] J. Shi and J. Malik. Normalized cuts and image segmentation. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 22(8):888–905, 2000.
- [29] M. Svensen. *GTM: The Generative Topographic Mapping*. PhD thesis, Aston University, Birmingham, UK, 1998.
- [30] P. Tino and I. Nabney. Hierarchical gtm: constructing localized non-linear projection manifolds in a principled way. (*IEEE*) *Transaction on Pattern Analysis and Machine Intelligence*, 24(5):639–656, 2001.
- [31] B. Zhang. Generalized k-harmonic means – boosting in unsupervised learning. Technical report, HP Laboratories, Palo Alto, October 2000.
- [32] B. Zhang, M. Hsu, and U. Dayal. K-harmonic means - a data clustering algorithm. Technical report, HP Laboratories, Palo Alto, October 1999.